

CHOICE IN A VARIABLE ENVIRONMENT: VISIT PATTERNS IN THE DYNAMICS OF CHOICE

WILLIAM M. BAUM AND MICHAEL DAVISON

THE UNIVERSITY OF AUCKLAND

Molar and molecular views of behavior imply different approaches to data analysis. The molecular view privileges moment-to-moment analyses, whereas the molar view supports analysis of more and less extended activities. In concurrent performance, the molar view supports study of both extended patterns of choice and more local patterns of visiting the choice alternatives. Analysis of the present data illustrated the usefulness of investigating order at various levels of extendedness. Seven different reinforcer ratios were presented within each session, without cues to identify them, and pigeons pecked at two response keys that delivered food on variable-interval schedules. Choice changed rapidly within components as reinforcers were delivered and, following each reinforcer, shifted toward the alternative that produced it. If several reinforcers were delivered consecutively by one alternative, choice favored that alternative, but shifted more slowly with each new reinforcer. A discontinuation of such a series of reinforcers by the delivery of a reinforcer by the other alternative resulted in a large shift of choice toward that alternative. These effects were illuminated by analysis of visits to the two alternatives. Changes in visit length occurred primarily in the first postreinforcer visit to the repeatedly reinforced alternative. All other visits tended to be brief and equal. Performance showed multiple signs of moving in the direction of a fix-and-sample pattern that characterized steady-state performance in earlier experiments with many sessions of maintaining each schedule pair. The analyses of extended and local patterns illustrate the flexibility of a molar view of behavior.

Key words: choice, dynamics, molar view, visit patterns, fix and sample, key peck, pigeons

Two different views of behavior exist within behavior analysis (Baum, 2002). In the older view, inherited from 19th-century associationism, behavior is seen as consisting of discrete events, usually called responses, occurring at moments in time. In this view, explanations rely on antecedent and consequent stimuli occurring in close temporal contiguity to the responses. Because more extended patterns of behavior are treated as built up out of sequences of discrete responses, much as a house may be built up out of bricks, this view may be called molecular, exploiting the connotation of discrete elements combining to form larger wholes. The newer view, first clearly articulated around 1970 (Baum, 1973; Baum & Rachlin, 1969; Herrnstein, 1970), sees behavior as consisting of temporally extended activities. In contrast with the molec-

ular view, explanations in this view rely on temporally extended relations between activities, contextual stimuli, and consequences. Activities are seen as patterns of behavioral allocation among parts that are themselves activities. Because activities are seen as necessarily extended and nested within one another, this view may be called molar, exploiting the connotation of extendedness of wholes.

The terms molecular and molar have a long history of being used in a variety of ways, creating potential for confusion. Some clarity may be achieved by restricting their use to refer only to the two views of behavior and the different explanations to which they lead. When talking about analysis, particularly analysis at different levels of temporal extendedness, confusion may be avoided by using terms like extended and local or long-term and short-term to refer to the different levels of temporal extendedness (Baum, 2002). Processes revealed at different levels of temporal extendedness also are best referred to as more and less local or extended.

The difference between these views is paradigmatic; it cannot be resolved by experiments or data (Baum, 2002). The two may be tested against one another only in terms of

This work was carried out at the University of Auckland and was supported by various grants to Michael Davison. We thank Mick Sibley for looking after the pigeons and the master and doctoral students who helped to conduct the experiment.

Reprints may be obtained from either author: Michael Davison, Psychology Department, The University of Auckland, Private Bag 92019, Auckland, New Zealand (e-mail: m.davison@auckland.ac.nz), or William M. Baum, 611 Mason Street, #504, San Francisco, California 94108 (e-mail: wbaum@sbcbglobal.net).

plausibility, elegance, and comprehensiveness. The molecular view has been applied to all sorts of phenomena: avoidance (Dinsmoor, 2001), choice (Hinson & Staddon, 1983), schedule effects (Killeen, Hall, Reilly, & Kettle, 2002), shaping (Donahoe & Palmer, 1994), and rule-governed behavior (Baum, 1992; Malott, 1992). Because of its reliance on contiguity for its accounts, however, it succeeds only by resorting to hypothetical events that stretch the limits of plausibility and elegance (Baum, 1992, 2003). One aim of the present paper is to demonstrate the superior elegance and comprehensiveness of the molecular view by showing its versatility in analyzing behavior at various levels of extendedness without resort to hypothetical events.

Choice, usually measured by the logarithm of the ratio of responding at two alternatives, constitutes an extended pattern of behavior, or an activity, that includes at least two parts: responding at one alternative and responding at the other. Because each of these parts is itself an activity, a more local analysis than choice analyzes responding at each alternative into its parts, in accord with the understanding that every activity is composed of parts that also are activities (Baum, 2002).

Choice in performance on concurrent schedules of reinforcement has usually been analyzed with the equation

$$\log\left(\frac{B_1}{B_2}\right) = s \log\left(\frac{r_1}{r_2}\right) + \log b, \quad (1)$$

where B_1 and B_2 are behavior allocations, measured in time or responses, to Alternatives 1 and 2, r_1 and r_2 are reinforcer rates obtained from Alternatives 1 and 2, b is a measure of bias toward one alternative or the other arising from factors other than r_1 and r_2 , and s is the sensitivity of the ratio of behavior to the reinforcer ratio. This equation describes a relation between an extended pattern of behavior (choice) and a distribution of reinforcers expressed as the reinforcer ratio. More extended patterns, including choice as a part, along with other activities, could be studied (Baum, 2002). More local patterns within choice also may be studied.

Examining visits or bouts of responding might reveal local patterns within the more extended pattern of behavioral allocation between the two reinforced alternatives. Al-

though defining a bout of behavior has been a challenging problem (Shull, Gaynor, & Grimes, 2001), in performance on concurrent schedules a definition may be readily at hand because the performance is punctuated by reinforcers and switches between alternatives (changeovers). A visit may be defined as a series of responses beginning with either a reinforcer or a changeover (Baum & Aparicio, 1999).

In the molecular view, a relation like that in Equation 1 is considered "derived" and only valid pending discovery of the moment-to-moment relations that would explain it (e.g., Catania, 1981). Analyses and theories of concurrent performance based on the molecular view focus, for example, on predicting at which alternative the next response will occur (e.g., Hinson & Staddon, 1983).

Examining dynamics within choice at a relatively local level, Davison and Baum (2000, 2002, 2003) found that, when various concurrent pairs of variable-interval (VI) schedules were presented briefly and without cues, choice from reinforcer to reinforcer was affected by the alternative from which each reinforcer came (i.e., its source). Following a series of reinforcers from one alternative, for example, when choice was strongly favoring that alternative, a reinforcer from the other alternative moved choice during the next interreinforcer interval substantially toward or even beyond indifference. This shift occurred even after as many as seven reinforcers in a row from one alternative. These experiments revealed regular relations between choice and various source sequences up to eight reinforcers in length. Landon, Davison, and El-liffe (2002, 2003a) found such dynamical relations also within steady-state choice. Bell and Baum (2002) found indications of similar dynamics within the foraging of a group of pigeons.

Davison and Baum (2002, 2003) also found that still more local regularities occurred between reinforcers. In moment-by-moment and response-by-response analyses, they observed that immediately following a reinforcer or a source sequence of reinforcers, a "pulse" of choice occurred, favoring the just-reinforced alternative at first and then shifting away with more time or responses, first rapidly and then ever more slowly, approaching an asymptote. Analyses in terms of visits

offered another approach, potentially more fruitful, but Davison and Baum (2003) only presented such an analysis briefly and without much comment.

Analysis of steady-state concurrent performance at a relatively extended level reveals patterns of visits. If schedule pairs are presented for a sufficient number of consecutive sessions, a pattern emerges in which the richer alternative receives the large majority of behavior, whereas the leaner alternative receives only occasional brief visits (Baum, 2002; Baum & Aparicio, 1999; Baum, Schwendiman, & Bell, 1999). Baum et al. (1999) called this pattern "fix and sample." It represents a local pattern even within visiting at the two alternatives and compares to predictions from optimal foraging theory (Houston & McNamara, 1981). Adopting a molar view of the sort of dynamical situation that Davison and Baum studied may reveal whether any tendency toward the fix-and-sample pattern occurs. It would appear in two regularities: (a) brief visits at the nonpreferred alternative that remain invariant across reinforcer ratios, and (b) a high probability of switching immediately after a reinforcer from the nonpreferred alternative.

The present experiment and analyses explored this possibility while examining the dynamics of choice as patterns of visits. As in the earlier studies, seven different pairs of VI schedules were presented within each session, each pair for a fixed number of reinforcers, either four or 12. Blackouts separated presentations of schedule pairs, and no cues accompanied the different pairs to distinguish them. In some conditions a changeover delay was in effect, in others, not. Patterns of behavior were examined following various source sequences of reinforcers.

METHOD

Subjects

Six experienced homing pigeons (Davison & Baum, 2000) numbered 91 to 96 were maintained at $85\% \pm 15$ g of their free-feeding body weights. Water and grit were available at all times. Designated body weights were maintained by weighing the subjects and feeding amounts of mixed grain immediately after the final session of the day.

Table 1

Sequence of experimental conditions, number of reinforcers per component, and the overall arranged rate of reinforcers in each condition of the experiment. In Conditions 10, 18, and 19, each component was presented twice per session.

Condition	Reinforcers per component	Reinforcers per minute	Changeover delay
10	4	6	2 s
11	12	6	2 s
16	12	6	None
17	12	2.22	None
18	4	2.22	None
19	4	6	None

Apparatus

The subjects were housed individually in cages (375 mm high by 370 mm deep by 370 mm wide) that also served as the experimental chambers. On one wall of the cage were three plastic pecking keys (20 mm diameter) set 100 mm apart center to center and 220 mm from a wooden perch situated 100 mm from the wall and 20 mm from the floor. Each key could be transilluminated by yellow, green, or red LEDs, and responses to illuminated keys exceeding about 0.1 N were counted as effective responses. Beneath the center key, and 60 mm from the perch, was a magazine aperture (40 mm by 40 mm). During reinforcement, the keylights were extinguished, the aperture was illuminated, and the hopper, containing wheat, was raised for 2.5 s. The subjects could see and hear pigeons in other experiments, but no personnel entered the room while the experiments were in progress.

Procedure

Sessions were conducted once per day and commenced at the same time each day. The 6 subjects were studied in succession with sessions lasting until a fixed number of reinforcers (see Table 1) had been collected, or until 45 min had elapsed, whichever occurred first. Sessions commenced with the left and right keylights illuminated yellow, which signaled the availability of a VI schedule on each key. Sessions were divided into seven components except in Conditions 10, 18, and 19, when 14 components were arranged. In those conditions, each of the basic seven components was

first selected randomly, and then a second randomization of these components was arranged. Each component lasted for a fixed number of reinforcers (see Table 1), and the components were separated by the blackout of both keys for 30 s. The arranged overall reinforcer rate was constant across components (see Table 1), but the values of the schedules on the two keys changed randomly over the components, providing seven different, unsignaled reinforcer ratios for the session. In all conditions, the reinforcer ratios in the seven components were 27:1, 9:1, 3:1, 1:1, 1:3, 1:9, and 1:27, and reinforcers were arranged dependently—that is, they were assigned randomly to one of the keys whenever a single VI schedule arranged a reinforcer. Sessions ended with the extinguishing of both keylights.

A changeover delay (COD; Herrnstein, 1961) was in effect only in Conditions 10 and 11. In those conditions, following a changeover to either key, a reinforcer could not be obtained for responding at the key switched to until 2 s had elapsed from the changeover (i.e., the first response at the key). Although data from Conditions 10 and 11 were treated in an earlier paper (Davison & Baum, 2002), they were included here for the purpose of assessing the effects of the COD. Conditions 12 through 15, which intervened between Conditions 11 and 16, also were treated in the earlier paper (Davison & Baum, 2002).

During sessions, the time of every event, coded by event type, was collected for detailed analysis. Training on each condition continued for 50 sessions, and the performances in the last 35 sessions were used in the data analyses.

RESULTS

To compare with previous results, the first analysis investigated how log response ratios changed as a function of successive reinforcers delivered in each of the seven components. To do this, all responses from the beginning of a component to the first reinforcer were pooled across all 35 (or 70) presentations of the component, then all were pooled from the first reinforcer to the second, then from the second to the third, and so on. Thus the data were organized reinforcer by reinforcer with performance mea-

sured prior to each reinforcer since the beginning of the component or since the previous reinforcer (i.e., not cumulated across successive reinforcers).

For both component lengths (four and 12 reinforcers per component) and both reinforcer rates (2.22 and 6 per minute), response ratios changed with increasing number of reinforcers delivered in accord with the component reinforcer ratio. Figure 1 shows sensitivity (s in Equation 1) calculated by fitting least-squares regression lines to log response ratios prior to each successive reinforcer versus arranged log reinforcer ratios for the group data. Equation 1 generally described the data well. Sensitivity increased progressively from near zero prior to the first reinforcer to 0.65 for six reinforcers per minute (rpm) with a COD, to 0.51 (6 rpm with no COD), and to 0.45 (2.2 rpm with no COD) prior to the 12th reinforcer. The difference between the diamonds and the squares in Figure 1 shows that the COD increased sensitivity substantially. The difference between the squares and the triangles shows a small but consistent increase in sensitivity for the higher overall reinforcer rate. The filled symbols show the sensitivities from the three conditions with four reinforcers per component. As reported previously (Davison & Baum, 2000), the difference in reinforcers per component had no discernible effect; each shorter curve in Figure 1 follows along the comparable longer curve without any systematic deviation. Because the results for the conditions with shorter components mimicked those for the longer components, further analyses focused on the more informative longer components.

Figure 2 shows the results of an analysis of response ratios in the interreinforcer intervals after various sequences of reinforcers. It began with the response ratio at the beginning of a component, up to the first reinforcer. After the first reinforcer, the alternative (left or right) that was reinforced first was designated Alternative P, and the other, not-first-reinforced was designated Alternative N. After the first reinforcer and before the second reinforcer, one response ratio was calculated; pecks at P divided by pecks at N. After the second reinforcer, and before the third reinforcer, two response ratios were calculated (both P/N)—one for each of the two




Fig. 1. Sensitivity to reinforcer ratio (s in Equation 1) as a function of the number of reinforcers that have been delivered in a component, calculated using programmed reinforcer ratios and grouped data. Filled symbols indicate conditions in which components ended after four reinforcers. Unfilled symbols indicate conditions in which components ended after 12 reinforcers.

sequences PN (i.e., left-right or right-left) or PP (i.e., left-left or right-right)—and so on, up to the eight possible sequences of four reinforcers. These sequences came from various components, although the frequency with which each sequence occurred varied across components (e.g., a series of several reinforcers on the left would rarely occur in the 1:27 component). Figure 2 shows this analysis for the grouped data. The results for Condition 16 (6 rpm; no COD) appear in both panels (squares). The effect of the COD may be seen in the top panel. The effect of overall reinforcer rate may be seen in the bottom panel. The result from each condition appears as a sort of branching tree. The log response ratio before the first reinforcer (zero on the x -axis) was close to zero, indicating no average preference at the beginning of the component, as might be expected. Each successive reinforcer produced a shift of preference (log response ratio) toward the alternative from which it came. When all four reinforcers occurred on the left or the right (PPPP), preference shifted progressively toward the reinforced alternative. Whenever a shift of reinforcer source occurred, a shift in preference followed it. As reinforcers continued for an alternative, the effect of each successive reinforcer decreased

as more reinforcers were delivered (see sequences PPPP, PNNN, and PPNN). The effects of three reinforcers in a row from alternative N sufficed to overwhelm the effect of the first reinforcer (sequence PNNN) in large part, but not completely, as we shall see below.

The greater spread of the tree for the COD condition (top panel) shows that the COD magnified the effects of reinforcer source. The greater spread for the higher reinforcer rate (bottom panel) indicates that preference shifted more with successive reinforcers in the higher reinforcer rate than the lower. Effects of reinforcers beyond the fourth (not shown) continued to decrease, but discriminating these effects was difficult because of the doubling of data points at each reinforcer delivery. Regardless of whether a COD was used, and regardless of reinforcer rate, the ordering of preference across the various four-reinforcer sequences was always the same. All sequences ending with a reinforcer from Alternative P ended with preference for P (positive log response ratio), and all sequences ending with a reinforcer from Alternative N ended with preference for N (negative log response ratio). Closest to indifference (zero log response ratio) was always the sequence PPPN.

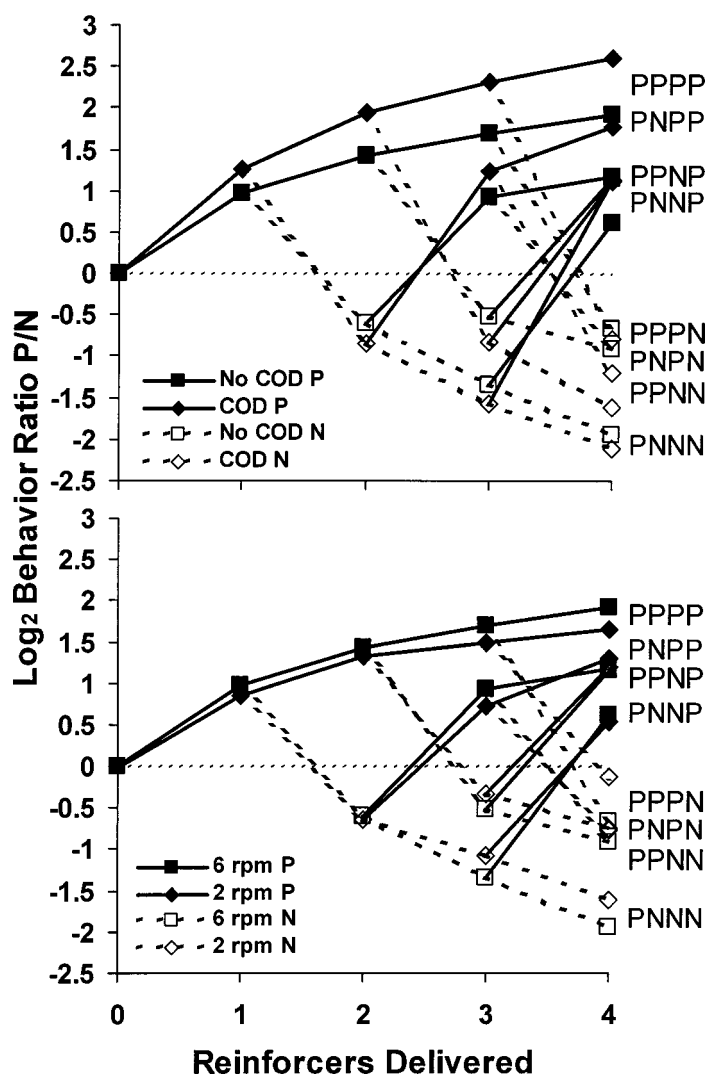


Fig. 2. Log interreinforcer peck ratio as a function of source sequence of reinforcers for the first four reinforcers within components. P: first-reinforced alternative. N: not-first-reinforced alternative. The letters to the right indicate the sequence of source alternatives (P and N). Filled symbols and solid lines indicate choice following a reinforcer from P. Unfilled symbols and broken lines indicate choice following a reinforcer from N.

Figure 3 shows, for the grouped data, effects of later reinforcers in a component for only part of the tree structure. It shows the effects of a sequence of left (or right) reinforcers, and also the effects of a single “discontinuation”—that is, a reinforcer from the right (or left)—at each sequential position, up to eight reinforcers. The results are shown for the same three conditions as in Figure 2. The COD magnified the shifts in response ratio. The higher reinforcer rate also resulted in larger shifts, but the effect was less than

that of the COD. Of 21 discontinuations, only two failed to switch preference to favor the alternative from which the discontinuing reinforcer came. Although the trend in the open symbols representing discontinuations is upward, toward indifference, the increasing lengths of the broken lines indicate that later discontinuing reinforcers shifted preference further, if only because continuing reinforcers shifted preference further from indifference.

In a previous paper, Davison and Baum




Fig. 3. Log interreinforcer peck ratio within components as a function of number of continuing reinforcers from the same alternative (filled symbols and solid lines) and following a reinforcer from the other alternative (i.e., a discontinuation; unfilled symbols and broken lines). P: first-reinforced alternative. N: not-first-reinforced alternative.

(2002) showed analyses that derived from the traditional molecular view of behavior, based on moments in time and momentary responses. With such an approach, one sees that the immediate effect of a reinforcer is to create a pulse of preference in favor of the source of the reinforcer. Figure 4 shows such pulses for the present grouped data and for Condition 11, which included the COD. Preference was calculated for the same sequences of continuing and discontinuing reinforcers as in Figure 3, but peck-by-peck, as P/N , in which P was the number of pecks at the just-reinforced alternative, and N was the number of pecks at the other, not-just-reinforced, alternative. Following the occurrence of one of the sequences, the location of the first peck—whether at the just-reinforced alternative or the other—was recorded, and all such first pecks were counted for the two alternatives. For the first peck (diamonds and squares), N also counted the number of postreinforcer

switches. The same was done for second pecks following one of the sequences, third pecks, and so on. Sample size necessarily decreased the later the peck after the first, because the next reinforcer terminated the counting of pecks. The integers on the x -axis show the number of reinforcers delivered in the component, as in Figure 3, so that the diamonds show the change in preference pulse with increasing numbers of continuing reinforcers from the beginning of a component, and the squares show the change in preference pulse following discontinuations after one or more continuing reinforcers at the first-reinforced alternative of the component. One square is missing from the top panel because, following a sequence of six continuing reinforcers plus a discontinuation, decreased postreinforcer switching due to the COD combined with small sample size (55 pecks) resulted in no first pecks at the not-just-reinforced alternative. The solid lines

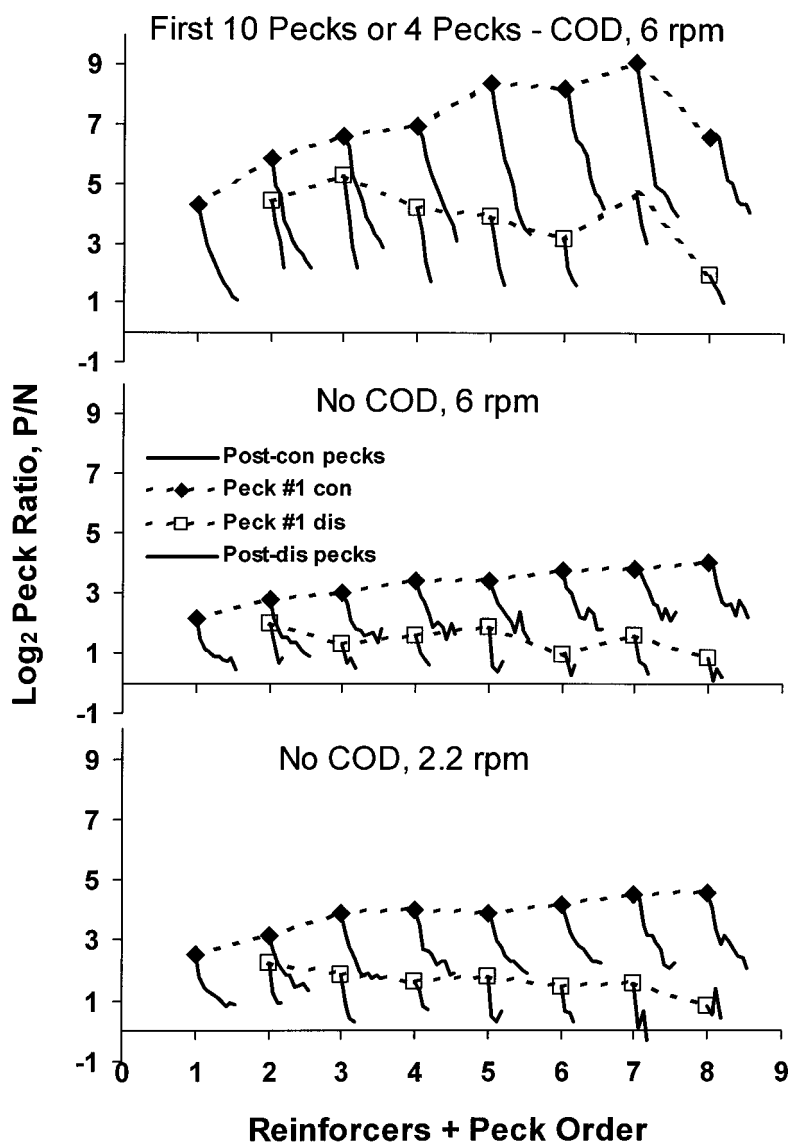


Fig. 4. Preference pulses following various numbers of continuing reinforcers from the same alternative (P; filled diamonds) and following a discontinuation of the continuing reinforcer series (a reinforcer from the other alternative, N; open squares) within components. Log peck ratio, peck by peck (tenths along the x -axis), is shown by the solid lines for the first 10 pecks following a series of continuing reinforcers and for the first four pecks following a discontinuation. Pulses increased in magnitude following more continuing reinforcers from the same alternative and decreased in magnitude following a discontinuation as the discontinuation was preceded by more continuing reinforcers.

show preference for pecks after the first—nine pecks for continuing sequences and three pecks for discontinuing sequences.

In all three conditions, the size of preference pulse tended to increase for continuing sequences of reinforcers and to decrease for discontinuations. The solid lines show that preference decreased toward indifference

with increasing temporal distance from the reinforcer. The COD resulted in larger pulses for continuing sequences—as high as 2^9 (i.e., 512:1)—in comparison with the two conditions without the COD. Pulses for continuing sequences were larger for the lower reinforcer rate than for the higher (log base 2 of 4.6 [equivalent to a ratio of 24:1]—vs.

4.1 [equivalent to 17:1]—after sequences of eight reinforcers in a row at the first-reinforced alternative). That the pulses following discontinuations (squares) were all smaller than pulses following continuing sequences indicates that continuing reinforcers had effects that lasted beyond the next reinforcer, just as shown by the more extended analysis of Figure 2. That the pulses following discontinuations also decreased as the number of continuing reinforcers increased indicates that the effects of continuing reinforcers also tended to cumulate, as may be seen in the more extended analysis of Figures 2 and 3.

Figures 1 to 4 show that the results of the conditions newly presented here resemble those reported previously for other conditions with the same pigeons (Davison & Baum, 2000, 2002). Figures 1 to 3 present analyses of extended patterns of behavior, in accord with the molar view of behavior, whereas Figure 4 shows a moment-by-moment analysis in accord with the molecular view of behavior. In the molar view, rather than leaping to moment-by-moment analysis, one may consider behavioral patterns at either more extended levels or at more local levels (Baum, 2002). What would be a more local pattern than the response ratio from one reinforcer to the next, on which Figures 1 to 3 are based? How do we examine patterns of behavior within the interreinforcer intervals?

Instead of moments or momentary responses, we may study bouts of activity or visits. Instead of the first peck following a reinforcer (Figure 4), Figure 5 shows the first visit to an alternative following a reinforcer, for various sequences of four reinforcers, as was done for interreinforcer response ratios in Figure 2. A graph is shown for each pigeon for Condition 17 (2.2 rpm with no COD). Visit length is measured as pecks per visit (PPV), calculated by summing pecks and visits across 35 sessions and dividing pecks by visits. A visit was defined as ending with a changeover; visits that lasted from one reinforcer to the next were omitted. Samples with fewer than five visits were omitted. As in Figure 2, P represents the first-reinforced alternative, and N represents the other, not-first-reinforced, alternative. The points at zero on the x -axis represent average visit length before any reinforcers were delivered, to give an indication

of bias as a difference in visit length. Four pigeons (91, 92, 94, and 95) showed noticeable bias.

By tracing sequences of diamonds or squares, one may see the effects of reinforcers when responding stayed at the just-reinforced alternative and when it switched postreinforcer to the other (hereafter called *stay* and *switch* postreinforcer visits, switch visits being relatively infrequent; see below). By comparing filled with open symbols, one may compare visits for the first-reinforced alternative (P) with those for the other (N). Two results may be seen readily in Figure 5. First, with continued reinforcers from the same alternative, the postreinforcer visit to that alternative increased in length. This may be seen in sequences of repeated filled diamonds and repeated open squares. For 2 pigeons (92 and 94), three reinforcers in a row from the not-first-reinforced alternative sufficed to increase the postreinforcer visit length up to be equal to that following four reinforcers in a row from the first-reinforced alternative. For the other pigeons, however, visit length increased to a lesser extent, retaining an effect of the first reinforcer. Second, postreinforcer visits at the not-just-reinforced alternative (much less frequent than visits at the just-reinforced alternative; see below) were always relatively short, often averaging two pecks and sometimes close to one (e.g., for Pigeon 96). Even following three continuing reinforcers, such switch postreinforcer visits after a discontinuation (unfilled diamond following three filled diamonds) were about as short as any (see Pigeons 91, 93, 94, and 95).

Figure 6 shows postreinforcer visits for sequences of four reinforcers calculated from the grouped data from Condition 17 (top panel) and the higher reinforcer rate (middle panel) and the condition with a COD (bottom panel). In all three conditions, repeated reinforcers from an alternative increased the postreinforcer visit length to that alternative. When the reinforcer rate was 6 rpm, three reinforcers in a row from the not-first-reinforced alternative (unfilled squares) sufficed to increase postreinforcer visit length to about the same level as following four reinforcers in a row from the first-reinforced alternative. Switch postreinforcer visits (un-

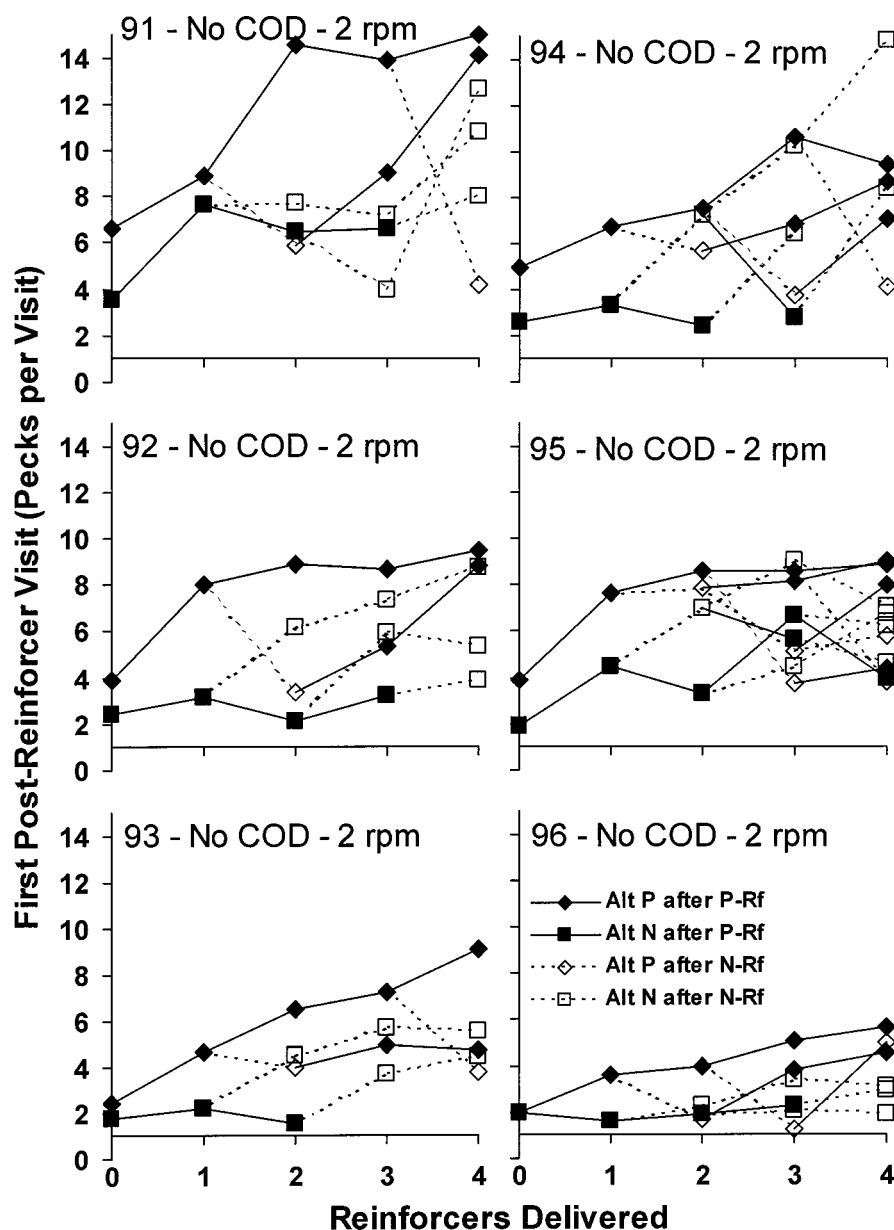


Fig. 5. Length of the first visit following a reinforcer (pecks per visit) for the condition with no COD and overall reinforcer rate of 2.2 per minute as a function of reinforcer-source sequence in the first four reinforcers; individual pigeons' data. P: first-reinforced alternative. N: not-first-reinforced alternative. Filled symbols represent switch (squares) and stay (diamonds) visit lengths following a P reinforcer. Unfilled symbols represent switch (diamonds) and stay (squares) visit lengths following an N reinforcer. Points plotted at zero on the x-axis show average visit length prior to the first reinforcer in a component. Missing points represent samples of fewer than five visits.

filled diamonds and filled squares) were always relatively short.

Comparing the two conditions with no COD, Figure 6 shows an effect of overall reinforcer rate in the relative spread of the two

tree-like graphs. Postreinforcer visits increased in length more for the lower reinforcer rate. Four continuing reinforcers resulted in an average postreinforcer visit length of 9.3 PPV for 2.2 rpm, in contrast




Fig. 6. Length of the first visit following a reinforcer (pecks per visit) as a function of reinforcer-source sequence in the first four reinforcers; grouped data. Top: condition with no COD and 2.2 rpm (same as in Figure 5). Middle: condition with no COD and 6 rpm. Bottom: condition with COD and 6 rpm. P: first-reinforced alternative. N: not-first-reinforced alternative. Filled symbols represent switch (squares) and stay (diamonds) visit lengths following a P reinforcer. Unfilled symbols represent switch (diamonds) and stay (squares) visit lengths following an N reinforcer. Points plotted at zero on the x-axis show average visit length prior to the first reinforcer in a component. Missing points represent samples of fewer than seven visits.

with 6.1 PPV for 6 rpm. Following three reinforcers in a row from the not-first-reinforced alternative (i.e., following the sequence PNNN), the average postreinforcer visit lengths differed less: 7.9 PPV for 2.2 rpm versus 7.0 PPV for 6 rpm. As might be expected from previous research, all visits were lengthened by the use of a COD. The bottom panel shows that even switch postreinforcer visits—when they occurred—never averaged fewer than 5.7 PPV. Even prior to any reinforcer delivery, visits at the nonpreferred alternative averaged 6.8 PPV (filled square at zero on the x-axis).

Figure 7 shows the effects on postreinforcer visits of sequences of continuing reinforcers followed by a discontinuation for individual pigeons in Condition 17 (no COD; 2.2 rpm). Missing points reflect sample sizes of fewer than five visits. As was true in Figure 5, continuing reinforcers increased the length of postreinforcer visits. Switch postreinforcer visits tended to be short and to decrease in length with increasing number of continuing reinforcers (except for Pigeon 96). A new result, not seen in Figure 5, is that postreinforcer visits to the not-first-reinforced alternative (unfilled squares) tended to decrease in length, becoming comparable to the switch visits after three or four continuing reinforcers (except for Pigeon 94).

These results are more clearly seen in Figure 8, which shows the postreinforcer visits for continuations and discontinuations calculated from the grouped data (sample sizes of seven or more visits). The top panel shows the group results based on the data in Figure 7. It clearly shows the downward trend in both switch and stay postreinforcer visit length following a discontinuation (unfilled diamonds and squares). Switch visits following a sequence of continuing reinforcers (filled squares) were comparably short or shorter. The middle and lower panels, for the higher reinforcer rate and the condition with the COD, show the same trends. In contrast with the lengthening of stay visits following continuing reinforcers (filled diamonds), stay visits following a discontinuation (unfilled squares) tended to shorten with increasing number of preceding continuing reinforcers. One exception was the lack of a downward trend in the postdiscontinuation switch postreinforcer visits (unfilled diamonds) for 6

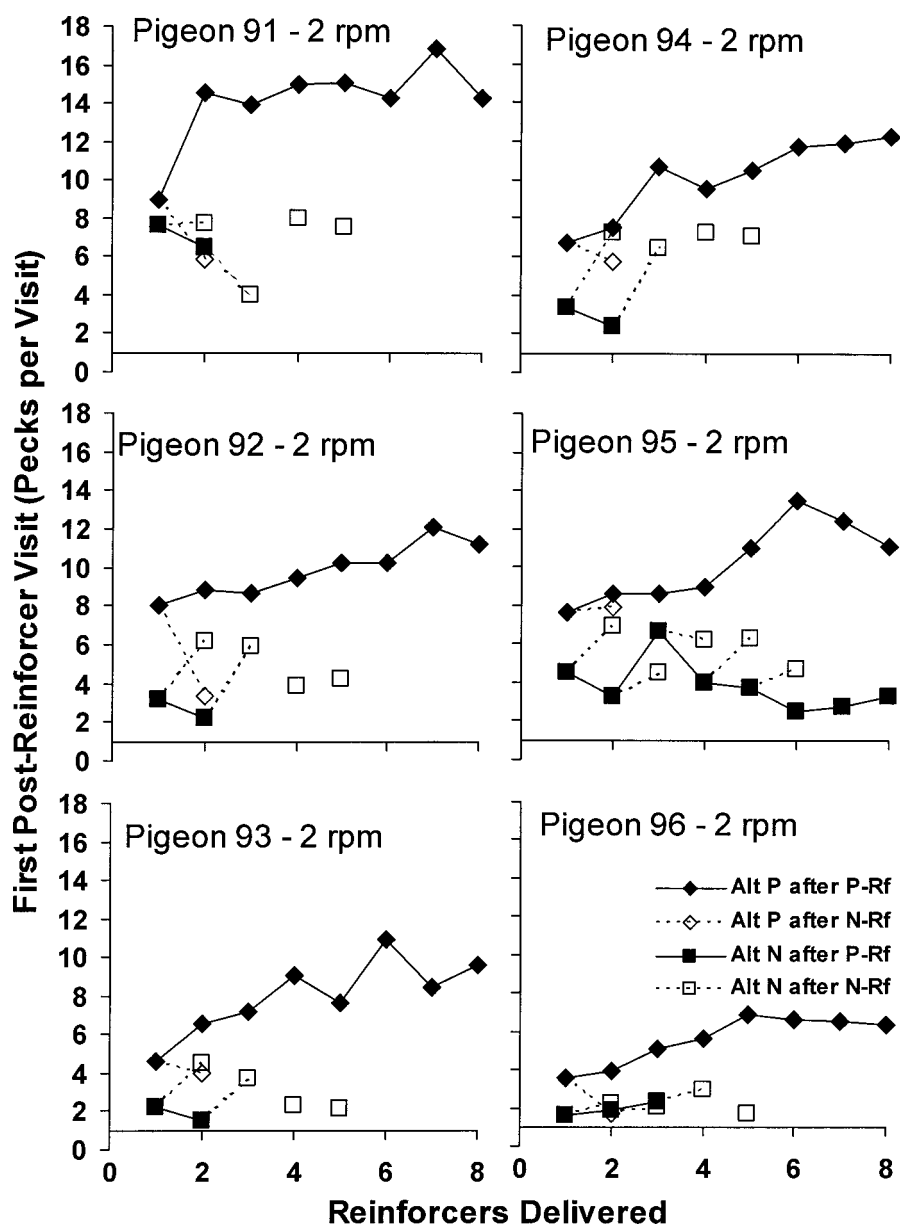


Fig. 7. Length of the first visit (pecks per visit) following a series of continuing reinforcers (filled symbols) or following a series of continuing reinforcers ending with a discontinuation (unfilled symbols) in the first eight reinforcers for the condition with no COD and overall reinforcer rate of 2.2 per minute; individual pigeons' data. P: first-reinforced alternative. N: not-first-reinforced alternative. Filled symbols represent switch (squares) and stay (diamonds) visit lengths following a series of P reinforcers. Unfilled symbols represent switch (diamonds) and stay (squares) visit lengths following a series of P reinforcers ending with an N reinforcer. Missing points represent samples of fewer than five visits.

rpm with no COD; these visits were uniformly short. Comparison of the top and middle panels reveals the same effect of overall reinforcer rate as in Figure 6. Postreinforcer visit length following seven or eight continuing

reinforcers reached 11.2 PPV for 2.2 rpm, but only 8.2 PPV for 6 rpm. Other types of postreinforcer visits also were shorter for the higher reinforcer rate, falling below 2 PPV for 6 rpm, but remaining above 2 PPV for 2.2 rpm. As

was seen in Figure 6, the COD lengthened all postreinforcer visits relative to the condition with no COD, even though the same trends were visible in the different types of visits. Even with a COD in use, stay visits following a discontinuation (unfilled squares) tended

to shorten with increasing number of continuing reinforcers.

Figures 5, 6, 7, and 8 focused on postreinforcer visits, omitting visits after the first visit following a reinforcer, but what about visits that followed a changeover rather than a reinforcer? Figure 9 shows, for the grouped data, visit lengths of the first 10 visits following a reinforcer in Conditions 16 (no COD; 6 rpm) and 17 (no COD; 2.2 rpm). The integers along the *x*-axis represent the number of reinforcers delivered, as in Figure 8, but between integers, at tenths, the next nine visits after the first postreinforcer visit are shown (lines without symbols). The points with symbols represent the same data as in Figure 8, except that, for purposes of comparison, visit length was calculated a bit differently. In Figure 8, postreinforcer visits that lasted all the way to the next reinforcer were excluded; in Figure 9, they were included. Comparison of the two figures reveals that whether such visits are included or not, the trends in the results are the same: Stay postreinforcer visits (filled diamonds in both figures) lengthened with increasing number of continuing reinforcers, but were calculated as shorter when reinforcer-to-reinforcer visits were included (Figure 9). The difference was larger for 6 rpm (8.2 vs. 7.0 PPV after eight continuing reinforcers) than for 2.2 rpm (11.2 vs. 10.8 PPV after seven continuing reinforcers). For the condition that included the COD (not shown), the difference was largest of all (22.9 vs. 11.8 PPV after seven continuing reinforcers), presumably because reinforcer-to-reinforcer visits tended to be much shorter than reinforcer-to-switch visits. Although all the same trends were visible with the COD as without, no re-




Fig. 8. Length of the first visit (pecks per visit) following a series of continuing reinforcers (filled symbols) or following a series of continuing reinforcers ending with a discontinuation (unfilled symbols) in the first eight reinforcers; grouped data. Top: condition with no COD and 2.2 rpm (same as in Figure 7). Middle: condition with no COD and 6 rpm. Bottom: condition with COD and 6 rpm. P: first-reinforced alternative. N: not-first-reinforced alternative. Filled symbols represent switch (squares) and stay (diamonds) visit lengths following a series of P reinforcers. Unfilled symbols represent switch (diamonds) and stay (squares) visit lengths following a series of P reinforcers ending with an N reinforcer. Missing points represent samples of fewer than seven visits.

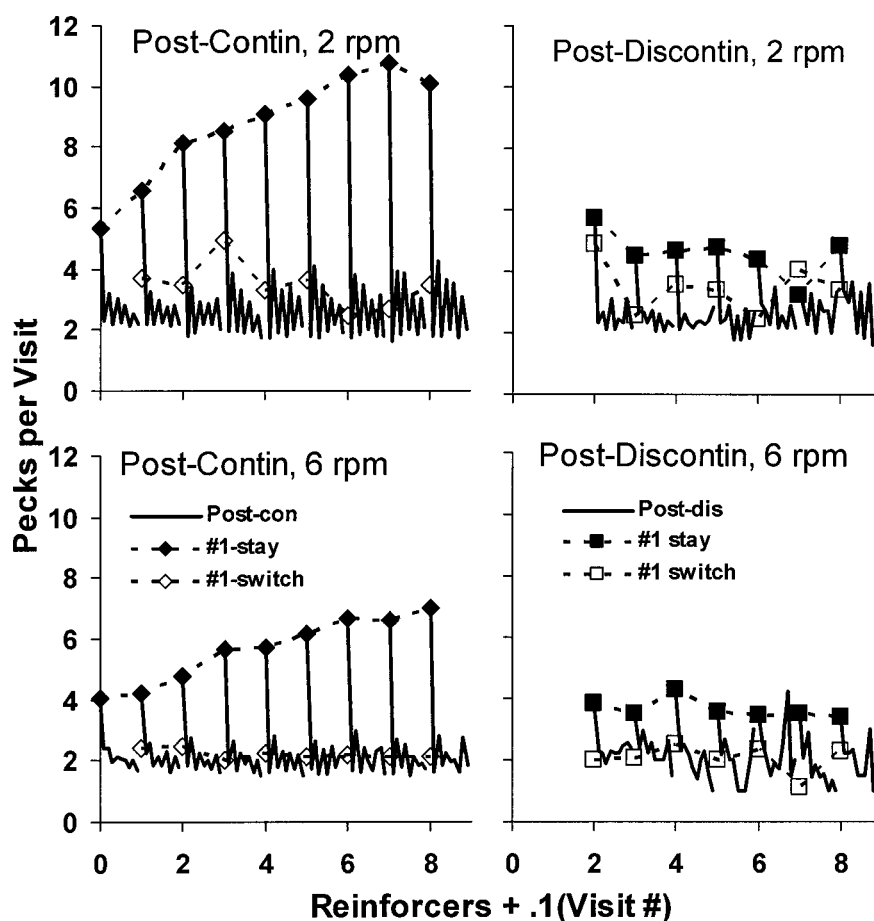


Fig. 9. Length of visits (pecks per visit) following a series of continuing reinforcers (diamonds) or following a series of continuing reinforcers ending with a discontinuation (squares) in the first eight reinforcers for the conditions with no COD and overall reinforcer rates of 2.2 and 6 per minute; grouped data. Filled symbols represent stay visits. Unfilled symbols represent switch visits. Graphs in the left column show visits following a series of continuing reinforcers. Graphs in the right column show visits following a series of continuing reinforcers ending with a discontinuation. Points without symbols represent postchangeover visit, plotted in tenths along the x-axis. Points plotted at zero on the x-axis show average visit length prior to the first reinforcer in a component. Missing points represent samples of fewer than seven visits.

sults from that condition appear in Figure 9 because so few postchangeover visits occurred that only a few averages could be calculated. Figures A1 to A6 in the Appendix show the same calculations for the individual pigeons, confirming that the group results were representative of the individual results.

The two panels on the left of Figure 9 show that, following a series of continuing reinforcers, visits after the postreinforcer visit (i.e., postchangeover visits) were invariably much shorter than the postreinforcer visit. The sawtooth, up-and-down pattern in the postchangeover visits, particularly apparent

with the lower reinforcer rate, shows a second-order residual effect of the continuing reinforcers on the postchangeover visits—that is, visits to the not-just-reinforced alternative continued to be shorter on average than to the just-reinforced alternative. This alternation of visit lengths also was observed in another experiment (Davison & Baum, 2003). It is clearly visible in the results for the individual pigeons as well (Figures A1 to A6). The difference diminished with increased number of switches because the postchangeover visits to the just-reinforced alternative shortened. After four continuing reinforcers,

switch postreinforcer visits (unfilled diamonds) became indistinguishable from other postchangeover visits.

The two panels on the right of Figure 9 show visits following a discontinuation. Although postchangeover visits (lines without symbols) tended to be shorter than postreinforcer visits (filled and unfilled squares), the difference was much less than following a series of continuing reinforcers because the postreinforcer visits following a discontinuation were much shorter. Figures A3 and A6 in the Appendix show that for 2 pigeons (93 and 96), the difference vanished. Indeed, the downward trend in the postdiscontinuation postreinforcer visits (filled and unfilled squares) brings them close to the postchangeover visits (see also Figures A2 to A6 in the Appendix). After a discontinuation following two continuing reinforcers, the switch postreinforcer visits (unfilled squares) became indistinguishable from other postchangeover visits. The absence of any clear sawtooth pattern in the postchangeover visits indicates the absence of any residual effect on the just-reinforced alternative.

Although Figures 7, 8, and 9 show the length of switch and stay postreinforcer visits, they give no indication of the relative frequency of these visits. Figure 10 shows the probability that pecking stayed at the just-reinforced alternative following a series of continuing reinforcers or a discontinuation. One graph appears for each pigeon in Condition 17 (no COD; 2.2 rpm). Missing points represent samples of fewer than nine visits. Two results stand out. First, the probability of staying following a series of continuing reinforcers was always high—usually higher than .9—and tended to increase with increasing number of continuing reinforcers. Second, the probability of staying after a discontinuation—that is, of staying at the just-reinforced not-first-reinforced alternative—although usually higher than .5, was almost always lower than the probability of staying after a continuation and tended to decrease with increasing number of preceding continuing reinforcers.

Figure 11 shows the probability of staying after a reinforcer calculated from the grouped data for the conditions with high and low reinforcer rate and with and without a COD. Missing points arose from sample sizes

of fewer than 10 visits. The triangles in the lower panel show the grouped data for the same condition as in Figure 10 (no COD; 2.2 rpm). The increase in reinforcer rate from 2.2 to 6 rpm lowered the probability of staying following a series of continuing reinforcers. This lowering also occurred following a discontinuation (unfilled triangles and diamonds) for shorter series of preceding continuations, but after six or seven continuing reinforcers, the probabilities of staying were about the same and substantially lower than the probability of staying after a continuing reinforcer (.64 vs .94 after seven continuing reinforcers). The top panel shows that the use of the COD increased the probability of staying after a series of continuing reinforcers to almost 1.0. Even so, an increase occurred for the first three continuing reinforcers. The unfilled squares show that, even with a COD, the probability of staying after a discontinuation decreased with the number of preceding continuing reinforcers. The one open square that appears to be missing is actually an anomalous probability of 1.0 coinciding with the sixth filled square—the same anomalous exclusive preference as resulted in a square missing from Figure 4.

To a good approximation, the main result shown in Figure 9 is that visits could be divided into two types: postreinforcer visits (filled symbols) and postchangeover visits (unfilled symbols and points without symbols). Such distinct types of visits occurring in order make up a pattern of behavior that is ignored in an analysis at a more extended level. For example, a typical analysis of choice between two alternatives pools responses at one alternative and at the other alternative without regard to any grouping of the responses into visits (e.g., the analyses summarized in Baum's, 1979, review). Recognizing the two types of visit, one may ask what effect sorting responses according to type would have on the more extended analysis, or, put differently, how ignoring the distinction of visits affects the extended analysis. One may ask how postreinforcer visits and postchangeover visits contribute to the results usually observed: sometimes matching (sensitivity equal to 1.0), but often undermatching (sensitivity less than 1.0).

Figure 12 shows the results of analyzing responding preceding the last four reinforcers

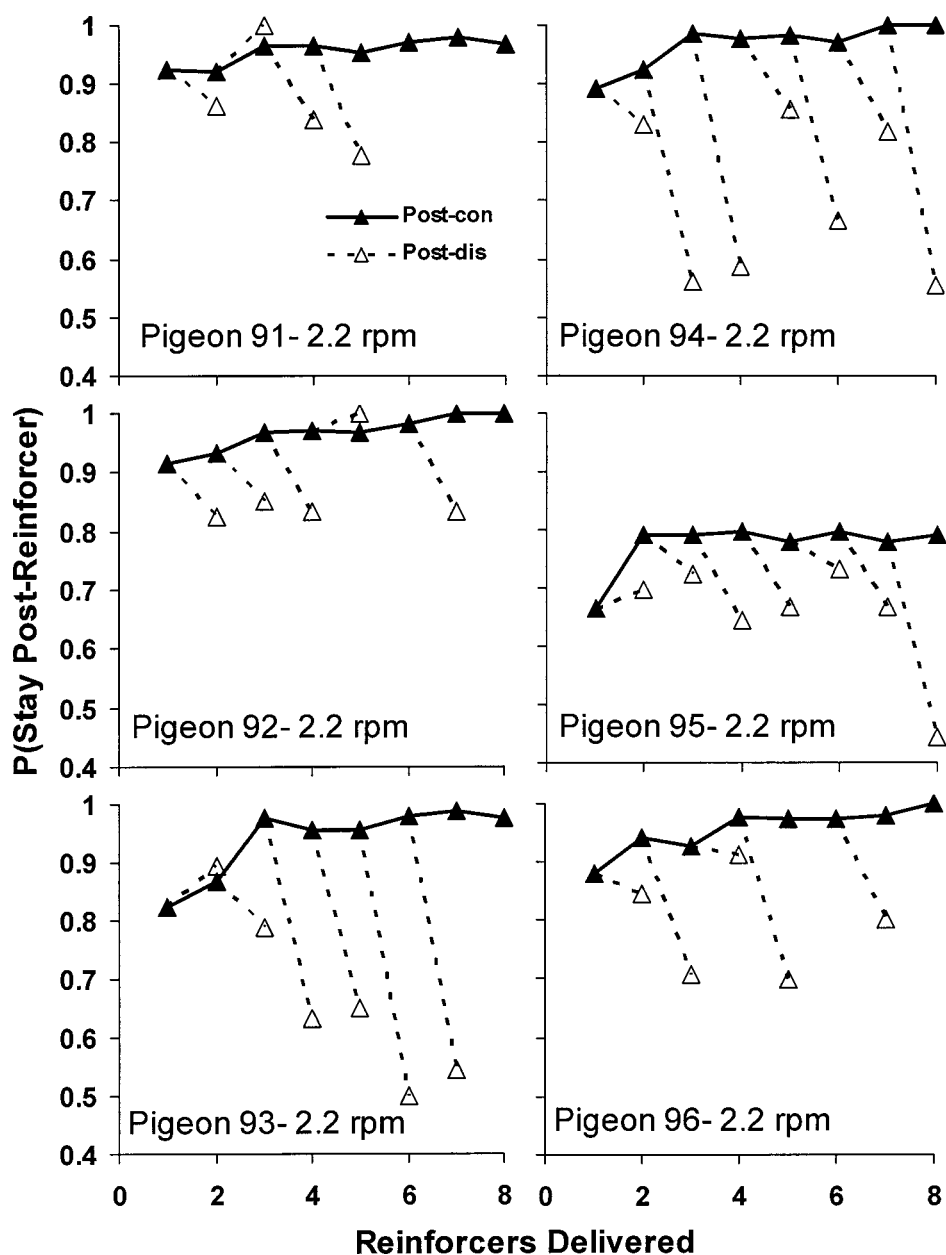


Fig. 10. Probability that responding stayed at the just-reinforced alternative following a series of continuing reinforcers (filled triangles) and following a series of continuing reinforcers ending with a discontinuation (unfilled triangles), for the condition with no COD and 2.2 rpm; individual pigeons' data. Missing points indicate samples of fewer than nine visits.

of components (i.e., from Reinforcer 8 to Reinforcer 12) in Condition 17 (no COD; 2.2 rpm) according to Equation 1. Sensitivity (Figure 1) appeared roughly stable across those last reinforcers of components. Each point in Figure 12 was calculated by aggregating

pecks at the left key and at the right key in a component, dividing left by right, and taking the logarithm (base 2) of the ratio. They are plotted against the logarithms (base 2) of the programmed reinforcer ratios. The lines were fitted to the three sets of




Fig. 11. Probability that responding stayed at the just-reinforced alternative following a series of continuing reinforcers (filled symbols) and following a series of continuing reinforcers ending with a discontinuation (unfilled symbols); grouped data. Top: comparison of the two conditions with and without a COD and overall reinforcer rate of 6 rpm. Bottom: comparison of the two conditions with 2.2 and 6 rpm and no COD.

points in each graph by the method of least squares. The equations of the three lines are shown in each graph. The triangles and middle equation represent the results of ignoring the distinction between postreinforcer and

postchangeover visits, as is usually done. The slopes (i.e., sensitivities), ranging from 0.29 to 0.59 across pigeons, are about what one would expect from Figure 1.

The diamonds and top equation represent the results of counting only pecks in stay postreinforcer visits (including reinforcer-to-reinforcer visits). The sensitivities all exceed 1.0, indicating overmatching, ranging from 1.4 to 1.8. This overmatching reflects the greater length of the postreinforcer visits to the richer alternative—that is, the alternative with more continuing reinforcers. The squares and bottom equation in each graph represent the results of counting only pecks in postchangeover visits (including switch postreinforcer visits and visits truncated by a reinforcer). The sensitivities, ranging from 0.05 to 0.25, are all lower than the corresponding sensitivities from the all-inclusive analysis (triangles and middle equations). Were the postchangeover visits equal at the two alternatives, the slopes of these bottom lines would equal zero. That the slopes are so small reflects the relative equality of the postchangeover visits, seen in Figure 9. Thus the undermatching reflected in the middle lines and Figure 1 may be seen as the outcome of combining the postreinforcer visits with the postchangeover visits. The greater the relative influence of the postchangeover visits, the more the undermatching. If the postchangeover visits became short and less frequent, less undermatching would occur. Table 2, which summarizes the results of the same analysis for Conditions 11 and 16, confirms this expectation: The slopes for postchangeover visits were closer to zero, and less undermatching occurred for the usual analysis ("Overall"). It also shows that the effect of the COD was to decrease the overmatching in the postreinforcer visits, presumably because the pecks made during the COD tended to equalize the postreinforcer visits to the two alternatives. (See Baum, 1974, for an analysis of responding during and after the COD.)

Whether the distinction between postreinforcer visits and postchangeover visits is ignored, as in Figures 1, 2, and 3, or whether it is made, as in Figures 9 and 12, choice within components clearly depends on the sequence in which reinforcers are obtained from the two alternatives (Figures 2, 3, 6, 8,

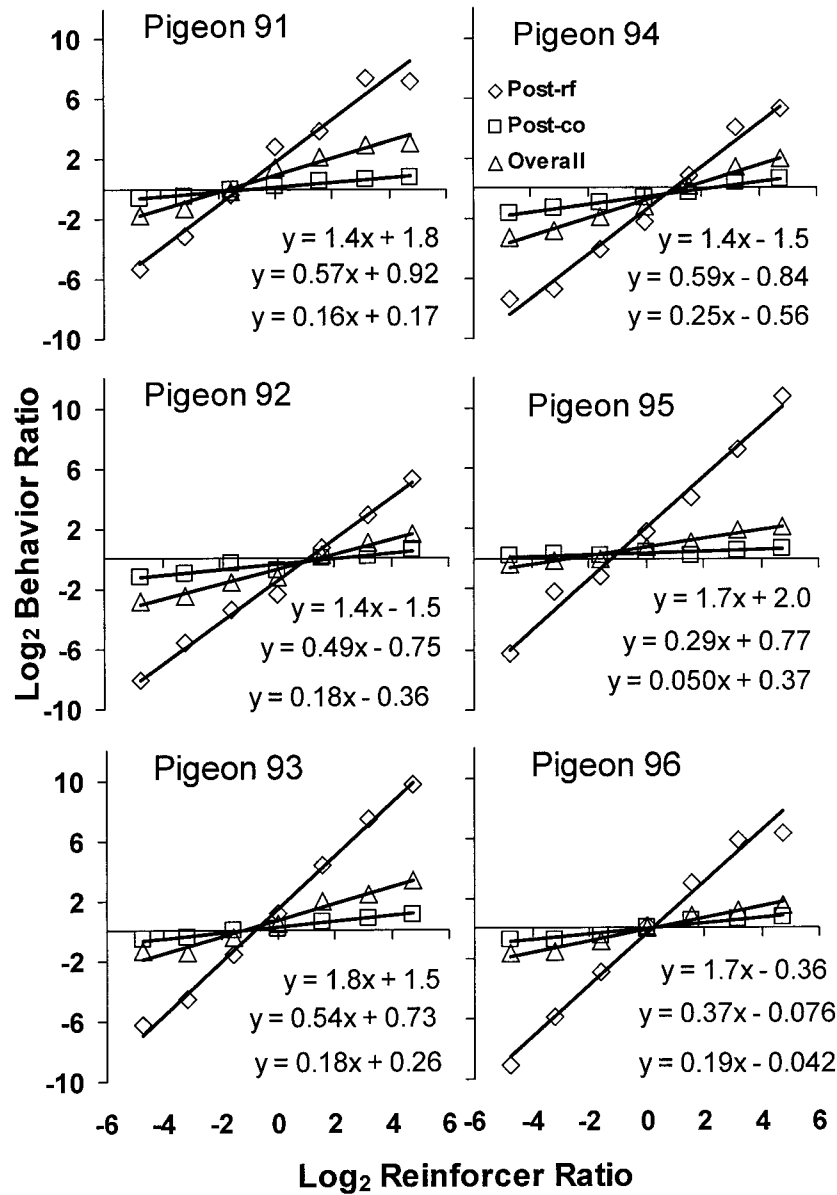


Fig. 12. Choice as a function of programmed reinforcer ratio across components, in accord with Equation 1 (the generalized matching law), calculated for all pecks (triangles), for pecks in visits preceded by a reinforcer (diamonds), and for pecks in visits preceded by a changeover (squares) from the eighth reinforcer of components to the 12th reinforcer, for the condition with no COD and 2.2 rpm; individual pigeons' data. Equations of the three regression lines appear for each pigeon.

and 9). For all these measures, we may ask how far back in the sequence the source of a reinforcer matters. We have shown the effects of different sequences up to eight reinforcers long, but have not asked thus far whether reinforcers eight or more back make any difference. For example, perhaps differences

following sequences of eight reinforcers were due only to the last six reinforcers. Thinking of the reinforcer sources as in a stack, we might phrase this question, "How deep was the control by past reinforcers?"

Figure 13 shows the difference due to the first reinforcer in a source sequence as a func-

Table 2

Analyses of the individual pigeons' data according to the generalized matching law for the last four reinforcers in components in Conditions 11 and 16 (6 rpm, with and without a COD), sorting pecks according to whether they occurred in a postreinforcer visit or in a postchangeover visit.

Pigeon	Postreinforcer		Postchangeover		Overall	
	Slope (<i>s</i>)	log(<i>b</i>)	Slope (<i>s</i>)	log(<i>b</i>)	Slope (<i>s</i>)	log(<i>b</i>)
No COD						
91	1.27	0.81	0.06	0.23	0.63	0.61
92	1.41	-1.42	0.05	-.30	0.60	-.67
93	1.41	0.89	0.08	0.14	0.64	0.65
94	1.44	-1.24	0.10	-.18	0.70	-.61
95	1.45	2.77	0.00	0.85	0.27	1.33
96	1.43	0.31	0.05	-.15	0.36	-.06
Group	1.31	0.06	0.03	0.16	0.50	0.21
COD						
91	1.21	0.30	0.01	-.07	0.69	-.09
92	1.04	-.37	-.01	-.10	0.63	-.14
93	1.17	0.88	0.00	0.69	0.66	0.96
94	1.07	-.22	-.07	-.18	0.54	-.15
95	1.12	-.75	-.03	-.95	0.69	-.65
96	1.09	-.14	-.06	0.58	0.59	0.36
Group	1.08	-.14	-.03	-.06	0.62	-.03

tion of the length of the sequence. Three curves are shown in each panel: one for log (base 2) peck ratio between reinforcers, as in Figures 2 and 3; one for log (base 2) postreinforcer visit length, as in Figures 5, 6, 7, and 8; and one for log (base 2) postchangeover visit length, as in Figure 9 (points without symbols). Each point represents the average difference in the measure following two source sequences, one beginning with a reinforcer from the first-reinforced alternative (P) and one beginning with a reinforcer from the not-first-reinforced alternative (N). For example, for sequences of four, the measure was calculated following the sequence PNNP and the sequence NNNP and the difference taken, following the sequence PPNN and the sequence NPNN and the difference taken, and so on. This was done for all sequences of each length. The frequencies of occurrence of the sequences depended on their length; ten sequences of two occurred within a 12-reinforcer component, nine sequences of three occurred, and so on. For sequences of two to five reinforcers, the differences were averaged across the ordinal positions of the sequences within components and across sequences. For sequences of six and seven reinforcers, adequate sample sizes required pooling across the ordinal positions and then

averaging the differences across sequences. Adequate sample sizes were unattainable for sequences of eight reinforcers.

Figure 13 shows these calculations for the grouped data of Conditions 11, 16, and 17. (Figures A7, A8, and A9 in the Appendix show that the group results were representative of the individual pigeons' results.) The graphs reveal that source sequences were no more than seven reinforcers deep—that is, the source of the seventh reinforcer back had a negligible effect. For log peck ratio (diamonds), the difference due to one reinforcer back (sequence length of two) was about 1.0 (i.e., a factor of 2). It declined to a minimum of 0.1—equivalent to a difference of 7%—for the condition with a COD (top panel), to a minimum of 0.18—equivalent to a difference of 13%—for 6 rpm with no COD (middle panel), and to a minimum of 0.05—equivalent to a difference of 3.5%—for 2 rpm (bottom panel). For postreinforcer visit length, the difference fell more rapidly with increasing number of reinforcers. It fell to a minimum of 0.04 at a depth of six reinforcers for the condition with a COD, to a minimum of 0.08 at a depth of five reinforcers for 6 rpm, and to a minimum of about zero at a depth of seven reinforcers for 2 rpm. For postchangeover visit length, the effect of previous




Fig. 13. The effect of the source of the first reinforcer in sequences of two to seven reinforcers on three inter-reinforcer measures following the sequence: \log_2 behavior ratio (diamonds); \log_2 first postreinforcer visit length (squares); \log_2 average postchangeover visit length at the not-first-reinforced alternative of the component (triangles); grouped data. The effect was calculated by taking the difference in a measure following two sequences identical except for the first reinforcer (first-reinforced alternative in a component vs. not-first-reinforced) and averaging across all sequences of the length on the x-axis.

reinforcer sources was generally small and, for the two conditions with 6 rpm (with and without a COD), was negligible for all sequence lengths. For the condition with 2 rpm, a small difference of about 0.2 (equivalent to a difference of 15%) appeared for sequences of two, three, and four reinforcers, but the influence of reinforcer sources five or more reinforcers back was negligible. For studying local effects of reinforcers, Figure 13 suggests that source sequences of six reinforcers in length might suffice. If this result proved reliable, for example, we could average across sequences of eight reinforcers (being conservative), effectively quadrupling our data in 12-reinforcer components. (See Landon et al., 2002, for a related analysis within standard long-term concurrent procedures.)

In the long term, Baum et al. (1999) found that concurrent performance eventually arrived at a pattern of fix and sample, in which staying at the richer alternative is interrupted only by occasional minimal visits to the leaner alternative. To what extent do the within-component changes in activity reflect this pattern, observed in typical long-term procedures in which conditions were presented for many sessions? If performance within components in the present experiment were moving toward fix and sample, we would expect to see changes in visit length and probability of switching. Visits to the more-reinforced alternative would grow longer (Figure 8), visits to the less-reinforced alternative would approach one peck in length (Figure 9), and the probability of staying after a reinforcer at the less-reinforced alternative would decrease toward zero (Figure 11). Of these three observations, the one least clearly supported by the analyses so far is the shortening of the visits to the less-reinforced alternative toward one peck, because Figure 9 suggests that the average visit length remained about two pecks. Analysis of the frequency distributions of visit length, instead of simply the means, might provide more understanding.

As Figures 8 and 9 show the means of visit lengths following sequences of continuing reinforcers ending either in a continuation or a discontinuation, Figure 14 shows, for Condition 17 (no COD; 2 rpm), the relative frequencies of different visit lengths following such sequences. Figures A10 to A15 in the Appendix show that these group results

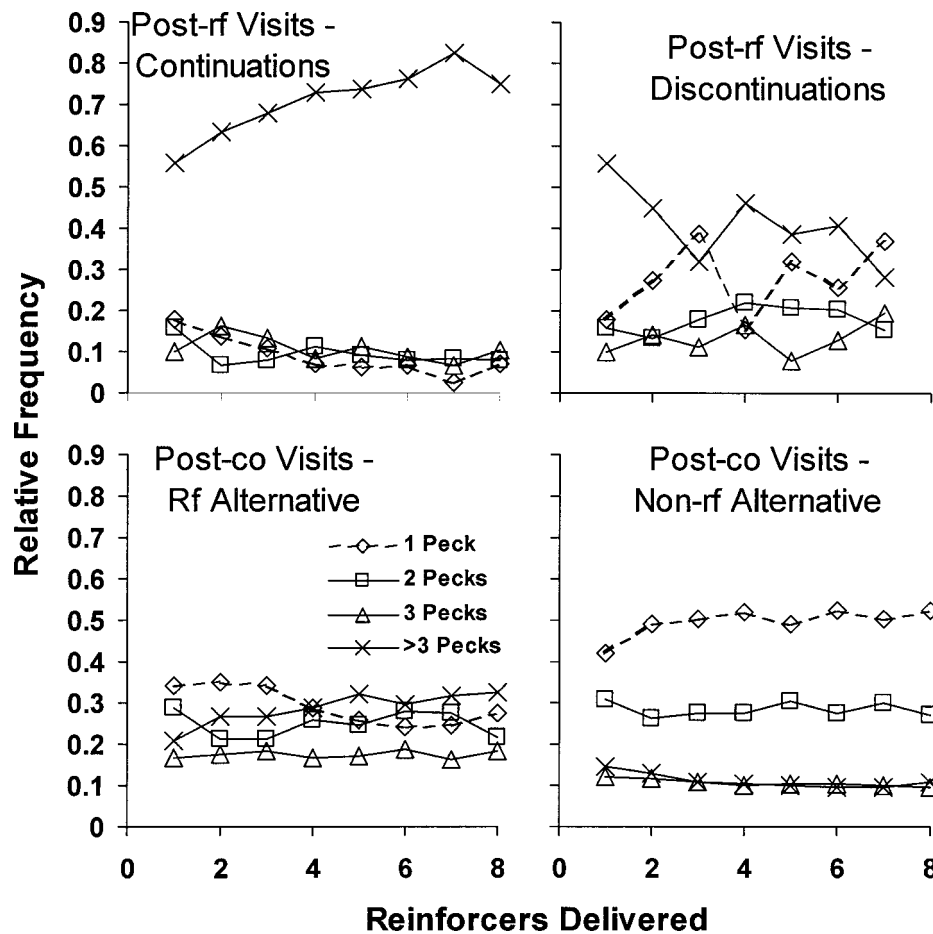


Fig. 14. Relative frequency of one-peck, two-peck, three-peck, and longer visits following delivery of the first one to eight reinforcers of components. Top row: first visit following a series of continuing reinforcers from the same alternative (left) or following a series of continuing reinforcers ending with a discontinuation (right). Bottom row: interreinforcer visits beginning with a changeover following a series of continuing reinforcers from the same alternative; left graph shows visits to the reinforced alternative; right graph shows visits to the nonreinforced alternative. Visits that ended with a reinforcer were excluded.

were representative of the individual pigeons' results. The two left panels and the top right panel show relative frequencies of visit lengths to the just-reinforced alternative; the lower right panel shows relative frequencies for the nonreinforced alternative following a series of continuing reinforcers. The top left panel shows that, following a series of continuing reinforcers, as more reinforcers were delivered, postreinforcer visits of more than three pecks increased in frequency, comprising 82% of visits after seven continuing reinforcers, whereas single-peck postreinforcer visits, although comprising only 18% of the visits even after one rein-

forcer, decreased in frequency, comprising 2.5% of visits after seven continuing reinforcers. In contrast, the top right panel shows that, following a discontinuation, the frequency of visits more than three pecks in length decreased with increasing number of preceding continuing reinforcers whereas the frequency of single-peck visits increased. Figure 11 shows also that the absolute frequency of all such visits (unfilled triangles) tended to decline. Were switches counted as zero-length visits, Figure 11 shows that the frequency of zero-length visits tended to increase, and a mean calculated with those included would have strengthened the de-

creasing trend shown in Figures 7 and 8 (unfilled squares).

The lower panels in Figure 14 show the relative frequencies of the lengths of visits beginning and ending with a changeover following a series of continuing reinforcers. The left panel shows relative frequencies for the reinforced alternative, which would correspond to the richer alternative in a long-term experiment and would be included in the fixing part of the fix-and-sample pattern. Those frequency distributions were relatively flat—all the frequencies were low—but still, as more reinforcers were delivered, the frequency of visits longer than three pecks increased whereas the frequency of single-peck visits decreased. In contrast, the right panel shows that, for visits to the nonreinforced alternative, which would correspond to the sampling part of the fix-and-sample pattern, frequency distributions were peaked, predominantly comprising one- and two-peck visits. Of these, one-peck visits were always more frequent and increased to make up 52% of the distribution after eight continuing reinforcers. Two-peck visits declined in frequency to make up 27% of the distribution after eight continuing reinforcers. Visits of more than two pecks decreased from 27% to 20% of the distribution.

Asking about the relation between the results presented here for brief exposures and the results for long-term exposures, one may take a converse approach and inquire whether any of the dynamics shown here appear in a long-term experiment. In the Baum et al. (1999) experiment, even though performance stabilized on a fix-and-sample pattern, a change of schedules might have resulted in changes in performance similar to those seen here. This possibility was tested by reanalyzing data from the first sessions of conditions in the Baum et al. experiment. Figure 15 shows, for 1 representative pigeon (B26) and three successive changes of schedules (with no COD), details of visit lengths during the first session of exposure to the new pairs of schedules. Visit lengths transformed by adding 1 to the length and taking the logarithm (base 2) are displayed in the order of their occurrence along the x -axis. Visits to the left are displayed as positive, visits to the right are displayed as negative, and although too many visits occurred for alternation to be visible,

visits to the left and right alternated. A diamond indicates the occurrence of a reinforcer. Extension of a line beyond a diamond constitutes a postreinforcer visit to the just-reinforced alternative. The larger symbols show medians calculated from the steady-state data; the ones to the left for the preceding schedule pair, and the ones to the right for the new schedule pair ultimately. Among the larger symbols, one pair of points—for the lean alternative—always shows the fix-and-sample pattern. For that alternative, a square (median postreinforcer visit length) is at zero, and a triangle (median postchangeover visit length) is at 1.0, representing single-peck visits or sampling. The other square–triangle pair shows medians for the rich alternative, always greater in length.

The top graph in Figure 15 shows a transition from an extreme reinforcer ratio (1:256) favoring the right key to a less extreme ratio (4:1) favoring the left key. The session commenced with long visits to the right and single-peck visits to the left, in keeping with the squares and triangles immediately to the left in the graph. A change in performance occurred after about 100 visits and 13 reinforcers, as visits to the right shortened, and visits to the left lengthened. After 500 visits, many visits to the right consisted of a single peck and the rest mainly of two pecks. Of five postreinforcer visits, one was zero. The session ended after 819 visits. The triangles and squares to the right of the graph show that near perfect fix and sample developed on the right ultimately, but the graph shows that an approximation to fix and sample occurred even within the first session.

The middle graph shows the first session of exposure to a new schedule pair favoring the right (1:8) following the preceding schedule pair (4:1). The session began with many one- and two-peck visits to the right key, in keeping with the previous fix-and-sample pattern. After about 600 visits and 26 reinforcers, visits began to lengthen on the right and to shorten on the left. After about 800 visits, visits on the right were about as long as they would be ultimately, and many one-peck visits were occurring on the left. The session ended after 1,109 visits. The bottom graph shows the first session of the next condition, which favored the left (64:1; “d” meaning dependent scheduling). After about 450 visits and 33 reinforcers, per-

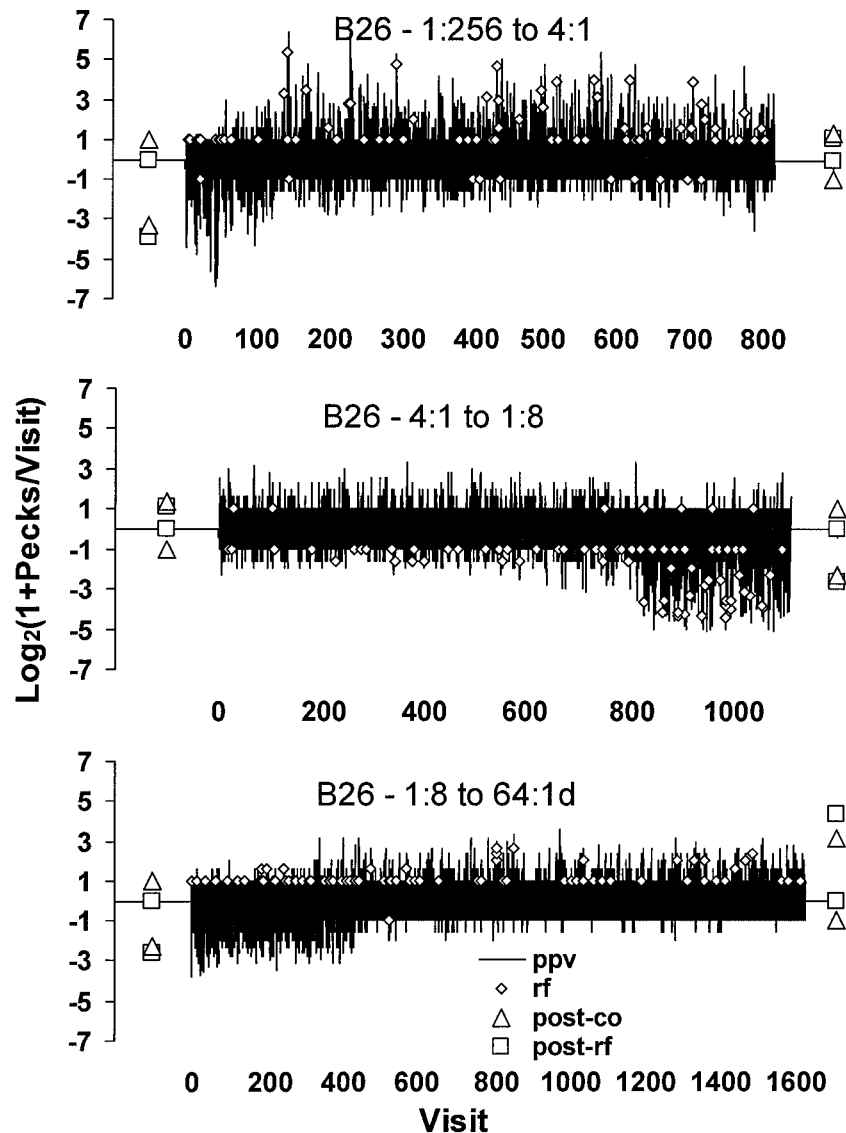


Fig. 15. Transition of choice within the first session of exposure to a new reinforcer ratio following many sessions of exposure to the preceding reinforcer ratio. Three successive transitions are shown for 1 pigeon in the experiment by Baum, Schwendiman, and Bell (1999). The y-axis represents the logarithm (base 2) of visit length transformed by adding 1.0. Visits are represented in order along the x-axis, visits to the left having positive ordinates, and visits to the right having negative ordinates. Open diamonds show the occurrences of reinforcers. Triangles and squares show median visit lengths for the steady-state performance in the preceding reinforcer ratio (to the left) and ultimately in the new reinforcer ratio (to the right). Squares show median length of the first postreinforcer visit (always zero at the lean alternative, because reinforcers were almost always followed by changeover to the rich alternative). Triangles show median length of visits beginning and ending with a changeover (always 1.0 at the lean alternative). A transition from a fix-and-sample pattern favoring the previously rich alternative to an approximation to fix and sample favoring the newly rich alternative occurred within the first session.

formance shifted from fixing on the right and sampling on the left to the reverse. After about 1,000 visits, visits to the right consisted predominantly of one peck. Visits to the left

lengthened at about the same time, but fell short of their ultimate length, shown by the upper triangle-square pair to the right. The session ended after 1,617 visits.

Two generalizations may be made about Figure 15. First, a shift toward the ultimate pattern of fix and sample occurred during the first session of exposure to a new schedule pair. Second, for these three transitions at least, the larger the change in reinforcer ratio, the quicker the shift in performance. If we think of the sessions shown in Figure 15 as containing long strings of continuing reinforcers with a few discontinuations and compare with Figure 7, the top panel of Figure 8, or the top panels of Figure 9 (recalling that the overall reinforcer rate in the Baum et al., 1999, experiment was 2 rpm), we see that the trends are the same: lengthening of the postreinforcer visits following continuing reinforcers and shortening of other visits.

DISCUSSION

Three main conclusions may be drawn from the results. First, the molar view of behavior provided a flexible approach to studying relatively local patterns of behavior in this dynamical situation. Second, the molar view provided a rationale for examining both the pattern of behavior across alternatives (i.e., choice) and the pattern of behavior within each alternative (i.e., visits). Third, the present results reveal a trend toward the same pattern as is seen in long-term experiments, namely, fix and sample (Baum et al., 1999).

Figures 1, 2, and 3 show the same local patterns of choice, from reinforcer to reinforcer, as appeared in earlier papers (Davison & Baum, 2000, 2002). Each reinforcer affected choice following it, continuing reinforcers shifted choice further and further toward the alternative from which they came, and shifts of reinforcement from one source to another (i.e., discontinuations) continued to have large effects on choice when examined at this level (Figures 2 and 3).

Figure 4 shows, with a molecular-inspired response-by-response analysis, that changes in choice occurred between reinforcers; that choice shifted immediately toward the alternative from which it came (the preference pulse), but then tended toward indifference, in keeping with the equality of reinforcer rates averaged across the whole session (Landon, Davison, & Elliffe, 2003a, 2003b).

An alternative to the moment-by-moment approach may be seen in the relatively local

analyses of visits shown in Figures 5 to 9. They help to explain the changes in choice seen in Figures 2 and 3. Figures 5 and 6, which examine postreinforcer visits following the first four reinforcers of a component, depending on source (i.e., the alternative from which a reinforcer was obtained), reveal shifts in postreinforcer visit length that underlay the shifts in choice from reinforcer to reinforcer seen in Figures 2 and 3. The correspondence between choice and postreinforcer visit length is most apparent if one compares the trends in the diamonds (visit length following continuing reinforcers and discontinuations) in Figure 6 with any of the tree-like graphs of Figure 2. Each behavior ratio shown in Figure 2 for a continuation, however, is composed of visits as parts. Each includes not only the postreinforcer visit on the first-reinforced alternative (P), but also subsequent visits to both the not-first-reinforced alternative (N) and to the first-reinforced alternative. Each behavior ratio for a discontinuation in Figure 2 similarly includes as parts both the postreinforcer visit (usually at N; Figures 10 and 11) and subsequent visits at both alternatives. Similar remarks apply to Figure 3 in comparison with Figure 8, the analyses of continuations and discontinuations across the first eight reinforcers of components. The diamonds in the three panels of Figure 8 may be compared to the three corresponding graphs in Figure 3.

The changes in visits at a relatively local level shown in Figures 6 and 8 may explain the shifts of choice at a more extended level shown in Figures 2 and 3. For continuing reinforcers for one alternative, the postreinforcer visit got longer and longer, making up the lion's share of the pecks averaged in with the other, shorter and more nearly equal, subsequent visits (Figures 9 and 12) to both alternatives when the behavior ratio was calculated. So the behavior ratio grew as the postreinforcer visit grew. For example, suppose that following the first reinforcer the postreinforcer visit contains six pecks (Figure 8, top panel), and subsequent visits contribute six pecks to the just-reinforced alternative and six pecks to the not-just-reinforced alternative (Figure 9); the behavior ratio would be 2.0 (\log_2 equal to 1.0; cf. triangles in Figure 3). Now say that following the eighth continuing reinforcer the postreinforcer visit contains 11 pecks (Figure 8, top panel) and

subsequent visits contribute five pecks to that alternative and four pecks to the other (Figure 9); the behavior ratio would be 4.0 (\log_2 equal to 2.0; cf. triangles in Figure 3). For discontinuations, Figures 6 and 8 show that postreinforcer visits (unfilled squares) were always shorter and grew shorter still with more preceding continuations. This explains the big shifts in choice seen in Figures 2 and 3 following discontinuations, because the shorter postreinforcer visit was averaged in with the short and nearly equal subsequent visits (Figure 9) when the behavior ratio was calculated. For example, suppose that the second reinforcer's source differs from the first reinforcer's, so it counts as a discontinuation, and the postreinforcer visit contains five pecks (Figure 8, top panel), and subsequent visits contribute five pecks to that alternative (N) and five pecks to the other (P; Figure 9); the behavior ratio (P/N) would be 0.5 (\log_2 equal to -1.0; cf. triangles in Figure 3). If the postreinforcer visit following a discontinuation preceded by seven continuations contained three pecks and the subsequent visits contributed six pecks to that alternative (N) and nine pecks to the other (P; see Figure 9, top right panel, in which visits to P are the longer ones), then the behavior ratio would be 1.0 (\log_2 equal to zero; cf. triangles in Figure 3).

The changes in visit length shown in Figures 8 and 9 may also be used to explain the preference pulses shown in Figure 4. The pulses result from the low probability of postreinforcer switches (Figures 10 and 11) combined with variability in postreinforcer visit length (Figure 14). When preference is calculated peck by peck, the low probability of a switch ensures that the great majority of first pecks occur at the just-reinforced alternative. For example, if the probability of staying at the just-reinforced alternative is .9 (Figure 11), the ratio of first pecks is 9.0 (\log_2 equal to about 3.2; diamonds in Figure 4). As the probability of staying at the just-reinforced alternative increases, so the calculated peck ratio increases also (Figure 4, filled diamonds). The probability of staying after a discontinuation was always less and decreased with preceding number of continuing reinforcers (Figure 11). Hence the peck ratios after discontinuations were smaller and decreased in Figure 4 (unfilled squares). The

peck ratio fell for the second and subsequent pecks, giving the steep decline that prompted the use of the word *pulse*, for two reasons. First, if a postreinforcer switch occurred, the second and perhaps third pecks were likely to be at the not-just-reinforced alternative (N). Second, although one-peck visits were relatively infrequent following continuations (diamonds in Figure 14), they did sometimes occur, and whenever they occurred, the second peck was at the not-just-reinforced alternative. Single-peck visits were more common following a discontinuation (Figure 14), resulting in a decrease in peck ratio for second pecks even there (Figure 4, unfilled squares). For the third and subsequent pecks, postchangeover visits increasingly affected the calculated peck ratios. One- and two-peck postchangeover visits were relatively common (Figure 14, lower panels), explaining why the peck ratios in Figure 4 declined further, eventually toward the sessional reinforcer ratio (indifference here, but see Davison and Baum, 2002, 2003, and Landon et al., 2003a, 2003b). Although explainable by the analysis of visits, the moment-to-moment results of Figure 4 give no inkling of the regularities in the patterns of visits revealed by the more extended analyses shown in all the figures subsequent to Figure 4.

As the behavior ratios shown in Figures 2 and 3 could be understood in the light of the changes in visit length, so, too, the growth of sensitivity shown in Figure 1 may be understood in that light. Figure 1 shows that sensitivity approached an asymptote as more reinforcers were delivered within components. For the eighth through 12th reinforcers, it was increasing only slowly. The slowness of the increase justified the analysis shown in Figure 12, which aggregated pecks within the last four reinforcers (excluding the 12th). Figure 12 revealed a large difference in sensitivity between postreinforcer and postchangeover visits. For postreinforcer visits, behavior ratios strongly overmatched reinforcer ratios; for postchangeover visits, behavior ratios strongly undermatched reinforcer ratios. The two sensitivities combined, however, to produce the intermediate undermatching shown in Figure 1. Because postchangeover visits changed little with delivery of reinforcers within components (Figures 9 and 14), the increasing sensitivity in Figure 1 resulted

primarily from increasing sensitivity in the postreinforcer visits. Sensitivity increased because postreinforcer visits to the rich alternative, which provided more continuing reinforcers, lengthened whereas postreinforcer visits to the lean alternative (represented by discontinuations) shortened (Figures 8 and 14).

The primacy of influence of the postreinforcer visits may explain the effects of adding a COD to a concurrent VI VI procedure (Herrnstein, 1961). Figure 1 shows what is well known: that a COD increases sensitivity, but never above 1.0 (e.g., Silberberg & Schrot, 1974). Herrnstein (1961), for example, reported that without a COD choice deviated relatively little from indifference as relative reinforcer rate varied, whereas with a COD choice tended to match relative reinforcer rate. Because Herrnstein exposed his pigeons to schedule pairs for too few sessions to allow development of the fix-and-sample pattern, without a COD visits to both alternatives would have been brief. Even if postreinforcer visits to the rich and lean alternatives differed, they would have been short, and many short and relatively equal postchangeover visits would have occurred. Such a performance would approximate simple alternation between alternatives, resulting in extreme undermatching. Adding the COD would have lengthened all visits, and particularly the postreinforcer visits. For example, an analysis of performance in Condition 11 (COD; 6 rpm) similar to that shown in the lower panels of Figure 9 for Condition 16 (no COD; 6 rpm) showed that all visits lengthened to eight or more pecks with a COD, but that the ratio of postreinforcer visit length to postchangeover visit length was reduced (about 11 to 8) in comparison with Figure 9, in which it equals about 6 to 2. If visits tended to be equal and long in Herrnstein's conditions with a COD, the number of pecks per reinforcer would have tended to be equal for the two alternatives. In the extreme, if every reinforcer occurred at the end of a COD, and eight pecks occurred during the COD, and no other pecks occurred, then eight pecks would accompany every reinforcer. The result would be perfect matching, because the ratio of pecks would always equal the ratio of reinforcers. For example, in a session of 60 reinforcers, if 40 came from one alternative and

20 from the other, 320 pecks would be made to the rich alternative and 160 pecks to the lean alternative. (See Houston, McNamara, & Sumida, 1987, for a similar discussion.)

Although the assumptions of eight pecks per COD and that all reinforcers occur at the end of a COD might be approximately correct, the assumption that all pecks occur during the COD is unrealistic, because, for example, pecks occur during postreinforcer visits, which entail no switch. In calculating choice, one would add the pecks made in the postreinforcer visits to the pecks made during the COD. If the postreinforcer visits at the rich alternative contained 11 pecks, and those at the lean alternative contained eight pecks, then the total pecks per reinforcer would be 19 and 16, producing a small degree of overmatching, but possibly indistinguishable from perfect matching. Whatever visits would occur that began and ended with a changeover would all be about equal at both alternatives (in our example, equal to eight pecks). They would contribute pecks equally to both alternatives, by themselves possessing zero sensitivity (s in Equation 1), and tending overall sensitivity toward undermatching. Supporting these speculations, Table 2 shows that the condition with a COD revealed slight overmatching ($s = 1.08$ for the group) for postreinforcer pecks, near zero sensitivity ($s = -.03$) for postchangeover pecks, and undermatching ($s = 0.62$) overall. The condition with 6 rpm and no COD, in contrast, showed strong overmatching ($s = 1.31$) for postreinforcer pecks, as in Figure 12, similar near zero sensitivity ($s = 0.03$) for postchangeover pecks, and undermatching ($s = 0.50$) overall.

This line of speculation might also explain the increase of sensitivity with higher reinforcer rate seen in Figure 1. The shorter interreinforcer intervals at 6 rpm allow less opportunity for the relatively equal postchangeover visits (Figure 9, lower panels), thereby reducing their influence on the calculation of overall choice. Carrying relatively more weight, the overmatching postreinforcer pecks tend to pull calculated sensitivity up for the higher reinforcer rate.

A study by Landon et al. (2003b), in which the reinforcer ratio across sessions differed from 1.0, produced evidence of both short-term dynamic effects like those shown here,

but also long-term effects of overall unequal reinforcer rates. Although all calculations of choice were shifted toward the long-term higher rate, the dynamic effects of reinforcer-source sequences resembled those reported here. Thus extended asymmetry in rate or magnitude of reinforcers (Davison & Baum, 2003) acts as an overlay on the local dynamics. Figure 13 indicates that all of these dynamic effects in visit length and choice occur within a time frame of just six or seven reinforcers. Analysis by Landon et al. (2002) indicated that a similar frame of no more than eight reinforcers applied to dynamics within steady-state performance. If this result holds generally, then analyses of dynamics in concurrent performance, both in short-exposure and long-exposure procedures, could use a moving window of seven or eight reinforcers. Such an approach would increase the data available for various source-of-reinforcer sequences and allow further analyses of performance as a function of varying sequences. One may ask, for example, how the dynamic effects interact with a more extended pattern of fix and sample.

The present analyses indicate that performance tended toward a fix-and-sample pattern (invariant minimal visits at the less-reinforced alternative), even if it failed to reach it within components. Figure 9 shows that visits to the less-reinforced alternative, both postreinforcer and postchangeover, tended to shorten with the delivery of more reinforcers. Figures 10 and 11 show that the probability of staying at the less-reinforced alternative following a reinforcer tended to decrease. Figure 14 shows that the frequency of one-peck visits tended to increase for the less-reinforced alternative. All of these trends might have continued if the procedure had allowed components to be longer. Time constraints and the need to present all seven components within a single session prevented components from being longer. Figure 15 shows, however, in an experiment with extended exposure, in which the fix-and-sample pattern developed, that a switch of schedules produced a new pattern of fix and sample within a single session but, particularly when the schedule change was relatively small, only after about 30 reinforcers (Baum et al., 1999). The possibility that longer components might

reveal further development of fix and sample remains to be explored in future research.

Even if concurrent-schedule performance always tends toward fix and sample, the logic of the molar view suggests that fix and sample, like any extended pattern, ought to be open to further analysis at a more local level. The reason is that, in the molar view, every activity is composed of less extended parts and is also a part of a more extended activity (Baum, 2002). Although analysis of the sample part of fix and sample might produce little profit, because sampling probably constitutes a minimal activity, analysis of the fix part of the pattern might offer helpful insights. When performance fixes on one alternative, all activities other than pecking are likely to occur as parts of the fixing pattern. Although these activities usually go unmeasured, both theory and observation suggest that taking account of them provides a better understanding of operant performance (Herrnstein, 1970; de Villiers, 1977). Thus an analysis of the activities that comprise fixing on the richer alternative should enhance our account of both steady-state and dynamic performance. This analysis also awaits future research.

REFERENCES

- Baum, W. M. (1973). The correlation-based law of effect. *Journal of the Experimental Analysis of Behavior*, 20, 137–153.
- Baum, W. M. (1974). Chained concurrent schedules: Reinforcement as situation transition. *Journal of the Experimental Analysis of Behavior*, 22, 91–101.
- Baum, W. M. (1979). Matching, undermatching, and overmatching in studies of choice. *Journal of the Experimental Analysis of Behavior*, 32, 269–281.
- Baum, W. M. (1992). For parsimony's sake: Comments on Malott's "A theory of rule-governed behaviour and organizational behavior management." *Journal of Organizational Behavior Management*, 12, 81–84.
- Baum, W. M. (2002). From molecular to molar: A paradigm shift in behavior analysis. *Journal of the Experimental Analysis of Behavior*, 78, 95–116.
- Baum, W. M. (2003). The molar view of behavior and its usefulness in behavior analysis. *Behavior Analyst Today*, 4, 78–81.
- Baum, W. M., & Aparicio, C. F. (1999). Optimality and concurrent variable-interval variable-ratio schedules. *Journal of the Experimental Analysis of Behavior*, 71, 75–89.
- Baum, W. M., & Rachlin, H. C. (1969). Choice as time allocation. *Journal of the Experimental Analysis of Behavior*, 12, 861–874.
- Baum, W. M., Schwendiman, J. W., & Bell, K. E. (1999). Choice, contingency discrimination, and foraging the-

- ory. *Journal of the Experimental Analysis of Behavior*, 71, 355–373.
- Bell, K. E., & Baum, W. M. (2002). Group foraging sensitivity to predictable and unpredictable changes in food distribution: Past experience or present circumstances? *Journal of the Experimental Analysis of Behavior*, 78, 179–194.
- Catania, A. C. (1981). The flight from experimental analysis. In C. M. Bradshaw, E. Szabadi, & C. F. Lowe (Eds.), *Quantification of steady-state operant behavior* (pp. 49–64). Amsterdam: Elsevier North-Holland.
- Davison, M., & Baum, W. M. (2000). Choice in a variable environment: Every reinforcer counts. *Journal of the Experimental Analysis of Behavior*, 74, 1–24.
- Davison, M., & Baum, W. M. (2002). Choice in a variable environment: Effects of blackout duration and extinction between components. *Journal of the Experimental Analysis of Behavior*, 77, 65–89.
- Davison, M., & Baum, W. M. (2003). Every reinforcer counts: Reinforcer magnitude and local preference. *Journal of the Experimental Analysis of Behavior*, 80, 95–129.
- de Villiers, P. (1977). Choice in concurrent schedules and a quantitative formulation of the law of effect. In W. K. Honig & J. E. R. Staddon (Eds.), *Handbook of operant behavior* (pp. 233–287). Englewood Cliffs, NJ: Prentice-Hall.
- Dinsmoor, J. A. (2001). Stimuli inevitably generated by behavior that avoids electric shock are inherently reinforcing. *Journal of the Experimental Analysis of Behavior*, 75, 311–333.
- Donahoe, J. W., & Palmer, D. C. (1994). *Learning and complex behavior*. Boston: Allyn & Bacon.
- Herrnstein, R. J. (1961). Relative and absolute strength of response as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behavior*, 4, 267–272.
- Herrnstein, R. J. (1970). On the law of effect. *Journal of the Experimental Analysis of Behavior*, 13, 243–266.
- Hinson, J. M., & Staddon, J. E. R. (1983). Matching, maximizing, and hill-climbing. *Journal of the Experimental Analysis of Behavior*, 40, 321–331.
- Houston, A. I., & McNamara, J. M. (1981). How to maximize reward rate on two variable-interval paradigms. *Journal of the Experimental Analysis of Behavior*, 35, 367–396.
- Houston, A. I., McNamara, J. M., & Sumida, B. H. (1987). Relative allocation on concurrent schedules can depend on schedule parameters when behavioral parameters are constant. *Journal of the Experimental Analysis of Behavior*, 47, 127–131.
- Killeen, P. R., Hall, S. S., Reilly, M. P., & Kettle, L. C. (2002). Molecular analyses of the principal components of response strength. *Journal of the Experimental Analysis of Behavior*, 78, 127–160.
- Landon, J., Davison, M., & Elliffe, D. (2002). Concurrent schedules: Short- and long-term effects of reinforcers. *Journal of the Experimental Analysis of Behavior*, 77, 257–271.
- Landon, J., Davison, M., & Elliffe, D. (2003a). Concurrent schedules: Reinforcer magnitude effects. *Journal of the Experimental Analysis of Behavior*, 79, 351–365.
- Landon, J., Davison, M., & Elliffe, D. (2003b). Choice in a variable environment: Effects of asymmetrical reinforcer distributions. *Journal of the Experimental Analysis of Behavior*, 80, 187–204.
- Malott, R. W. (1992). A theory of rule-governed behavior and organizational behavior management. *Journal of Organizational Behavior Management*, 12, 45–65.
- Shull, R. L., Gaynor, S. T., & Grimes, J. A. (2001). Response rate viewed as engagement bouts: Effects of relative reinforcement and schedule type. *Journal of the Experimental Analysis of Behavior*, 75, 247–274.
- Silberberg, A., & Schrot, J. (1974). A yoked-chamber comparison of concurrent and multiple schedules: The relationship between component duration and responding. *Journal of the Experimental Analysis of Behavior*, 22, 21–30.

Received February 23, 2003

Final acceptance: December 13, 2003

APPENDIX

FIGURES SHOWING CALCULATIONS FOR INDIVIDUAL PIGEONS

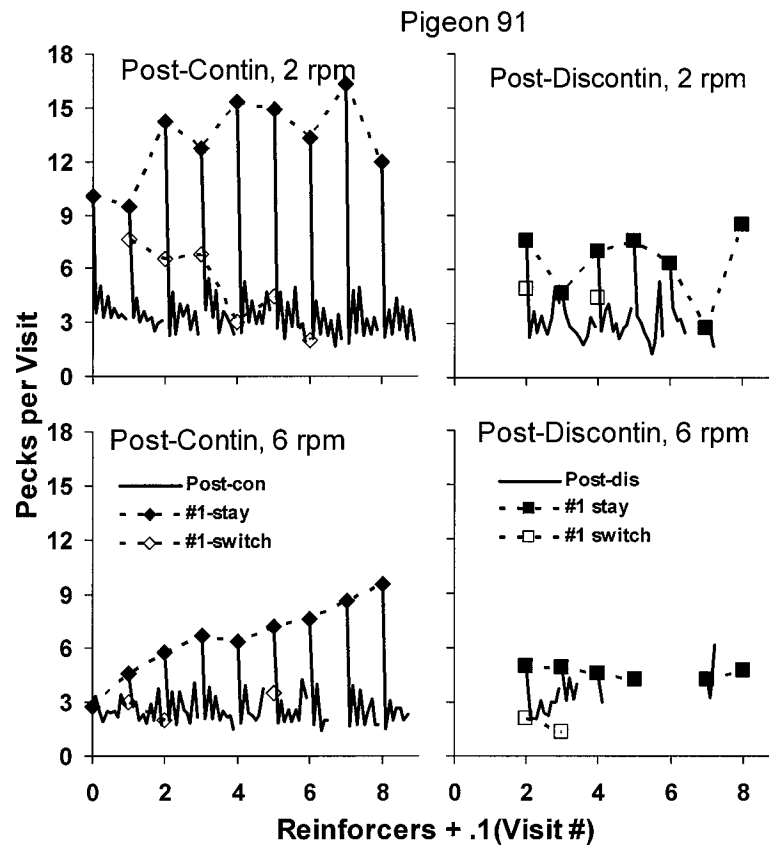


Fig. A1. Same analysis as shown in Figure 9 for Pigeon 91 (except missing points represent samples of fewer than three visits). See caption for Figure 9.

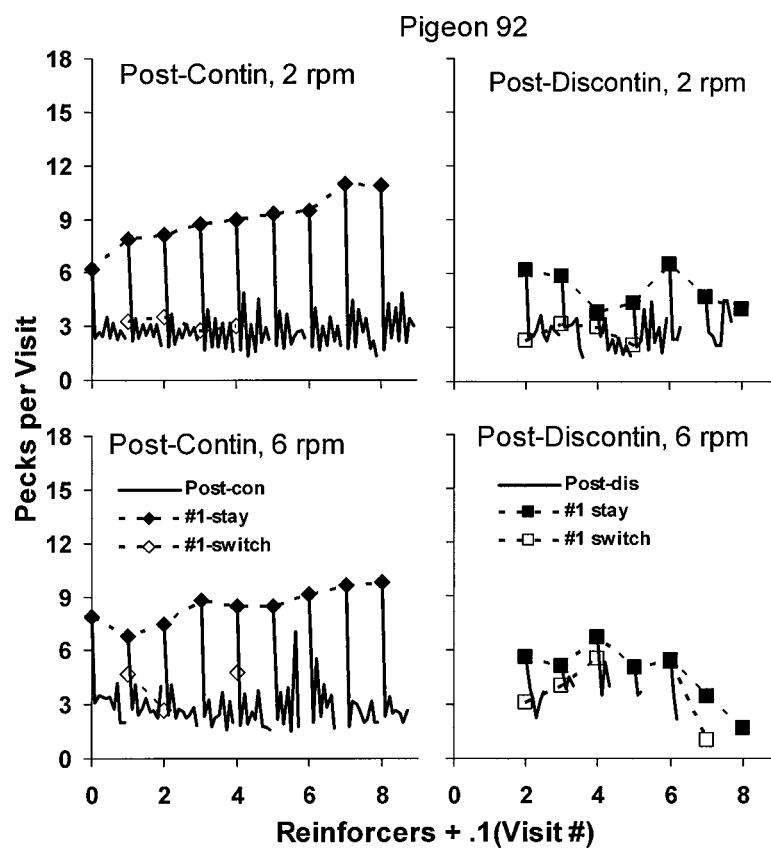


Fig. A2. Same analysis as shown in Figure 9 for Pigeon 92 (except missing points represent samples of fewer than three visits). See caption for Figure 9.

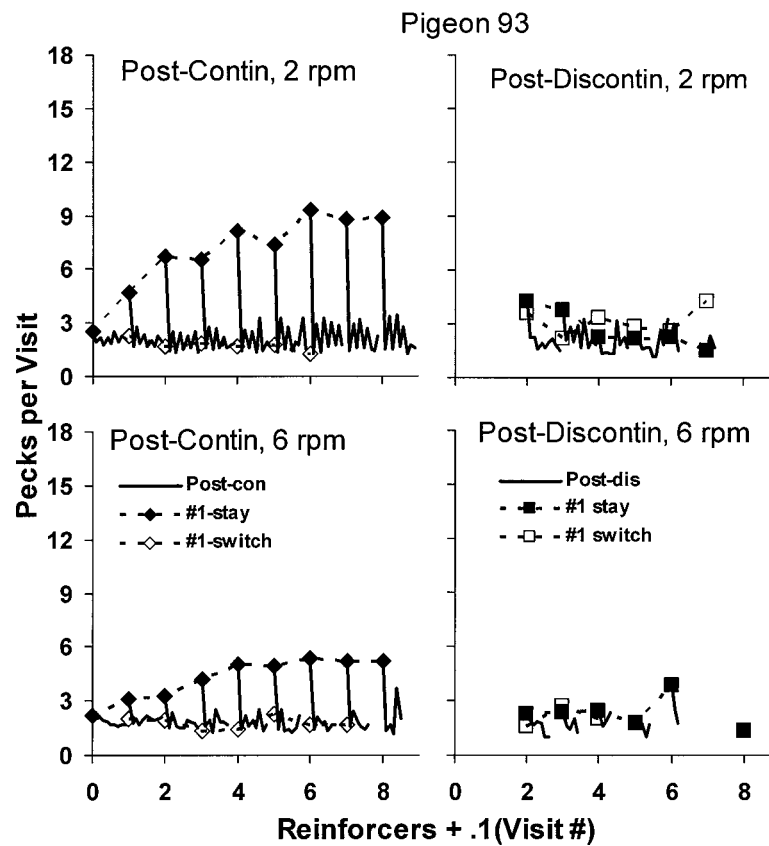


Fig. A3. Same analysis as shown in Figure 9 for Pigeon 93 (except missing points represent samples of fewer than three visits). See caption for Figure 9.

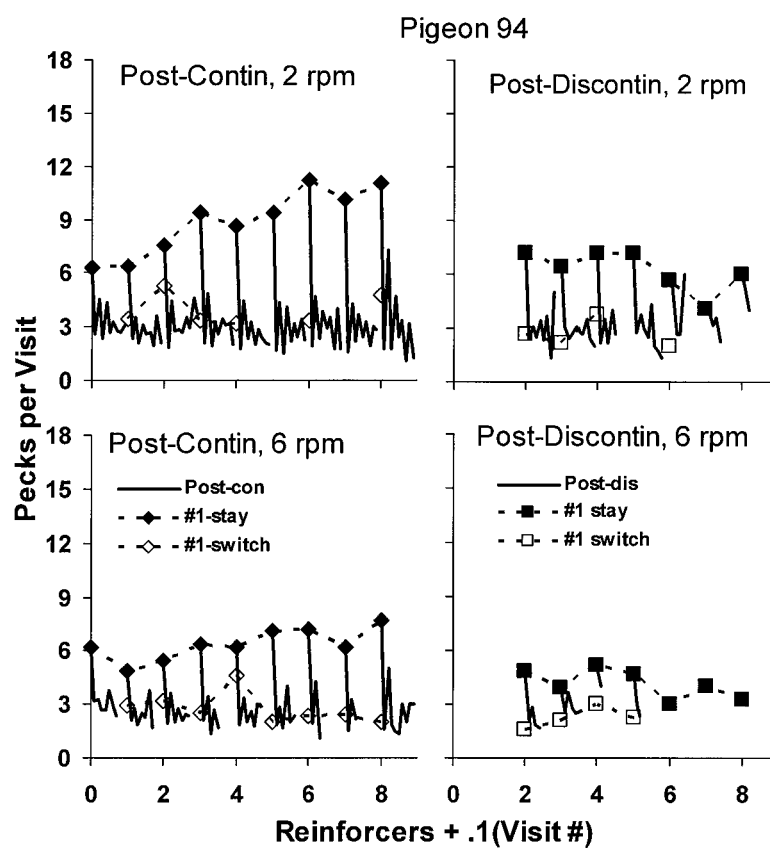


Fig. A4. Same analysis as shown in Figure 9 for Pigeon 94 (except missing points represent samples of fewer than three visits). See caption for Figure 9.

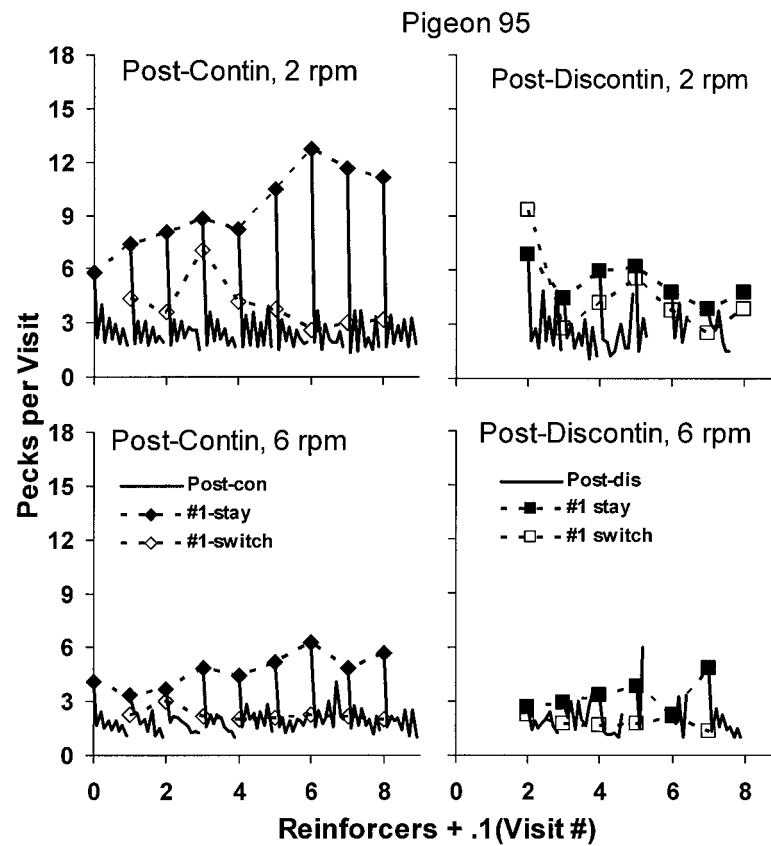


Fig. A5. Same analysis as shown in Figure 9 for Pigeon 95 (except missing points represent samples of fewer than three visits). See caption for Figure 9.

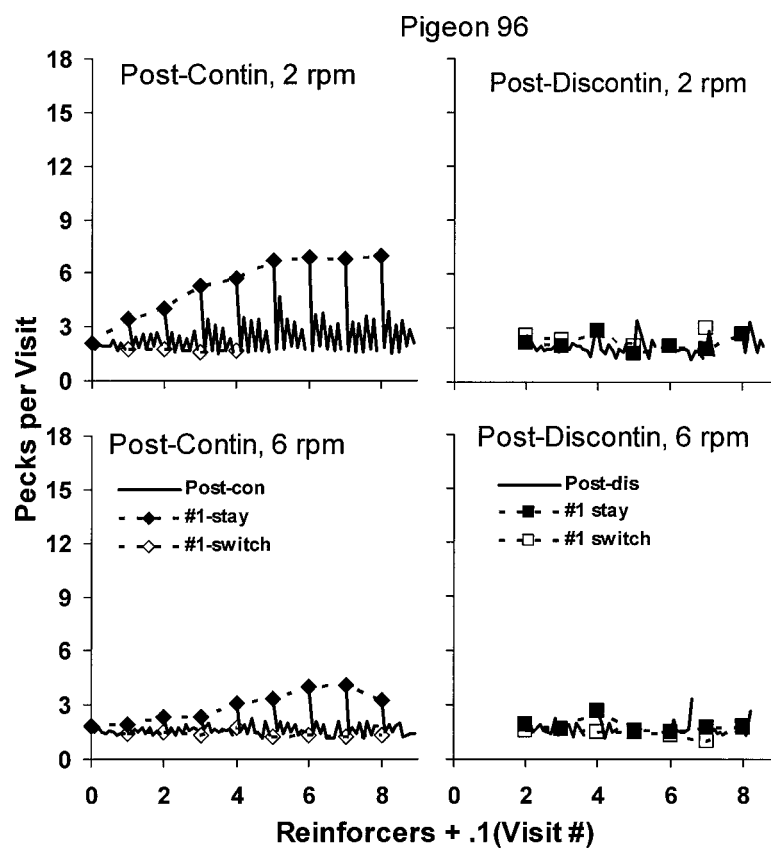


Fig. A6. Same analysis as shown in Figure 9 for Pigeon 96 (except missing points represent samples of fewer than three visits). See caption for Figure 9.

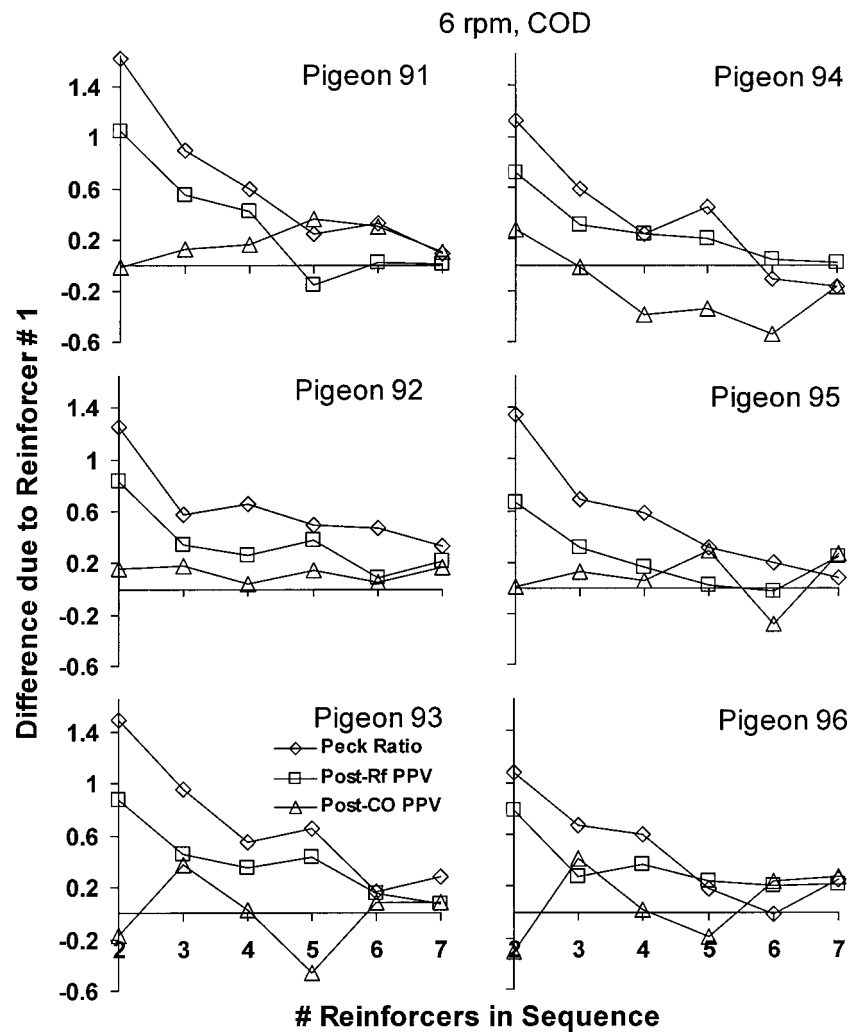


Fig. A7. Same analysis as shown in Figure 13 for individual pigeons in Condition 11. See caption for Figure 13.

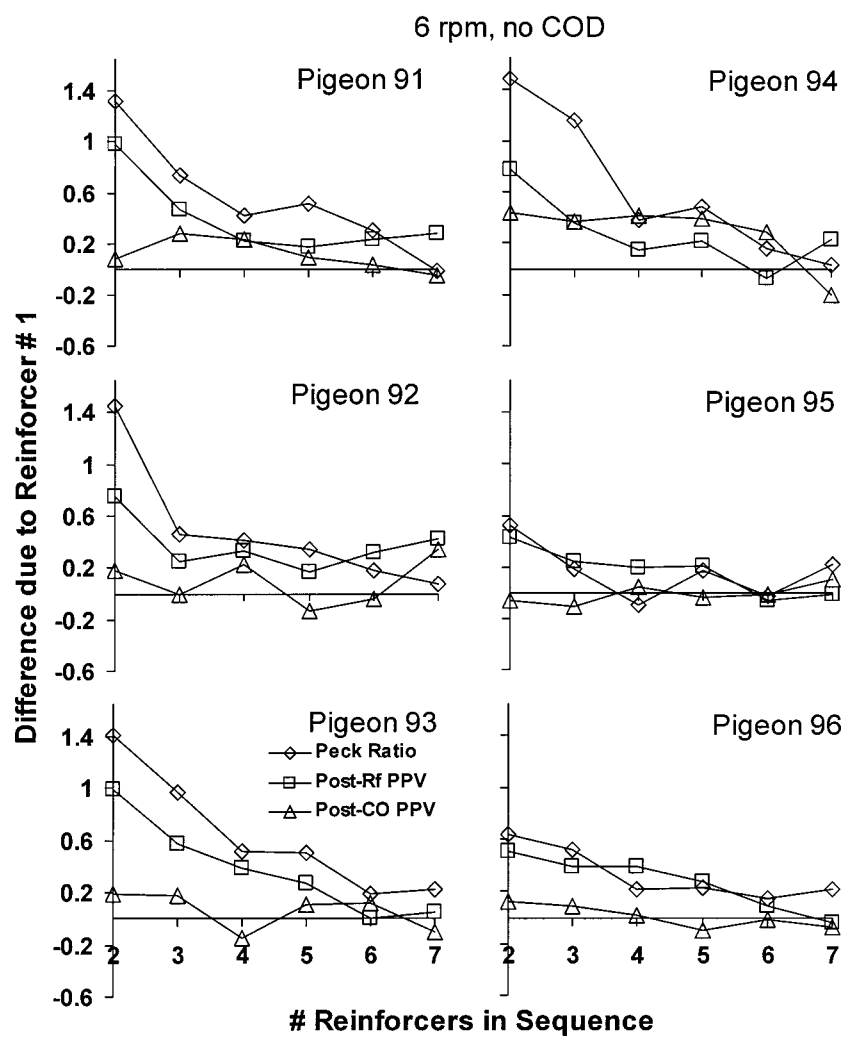


Fig. A8. Same analysis as shown in Figure 13 for individual pigeons in Condition 16. See caption for Figure 13.

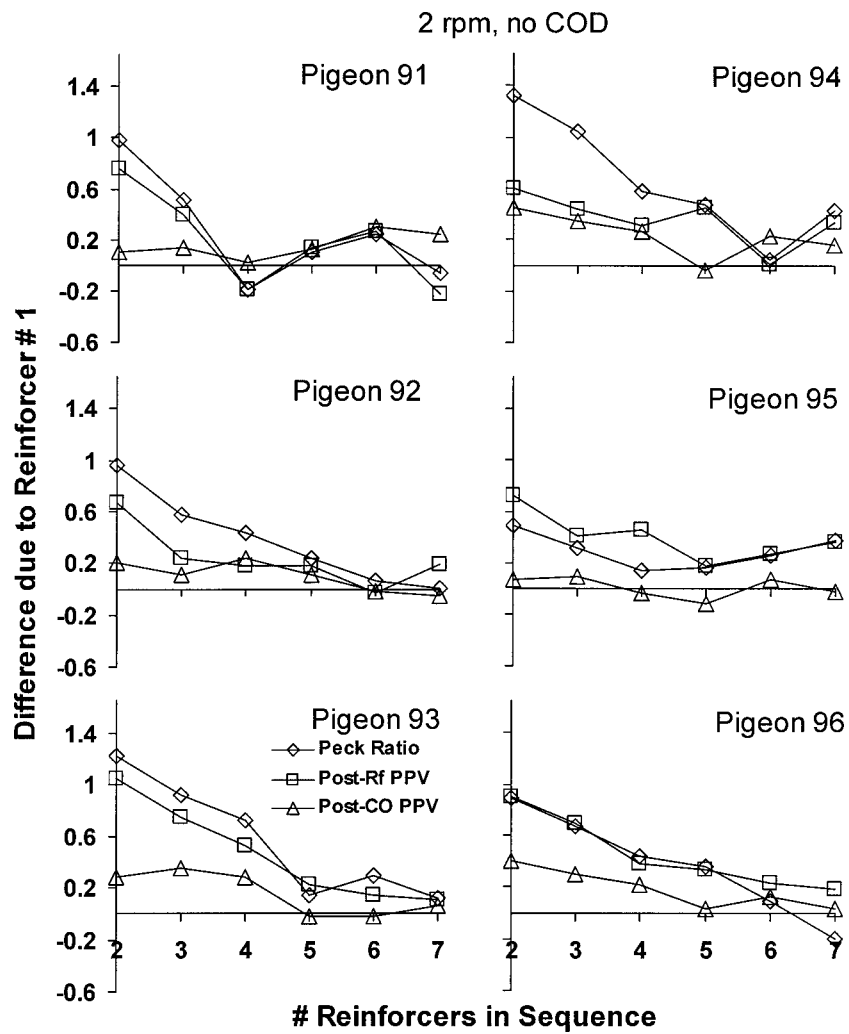


Fig. A9. Same analysis as shown in Figure 13 for individual pigeons in Condition 17. See caption for Figure 13.




Fig. A10. Same analysis as shown in Figure 14 for Pigeon 91. See caption for Figure 14.

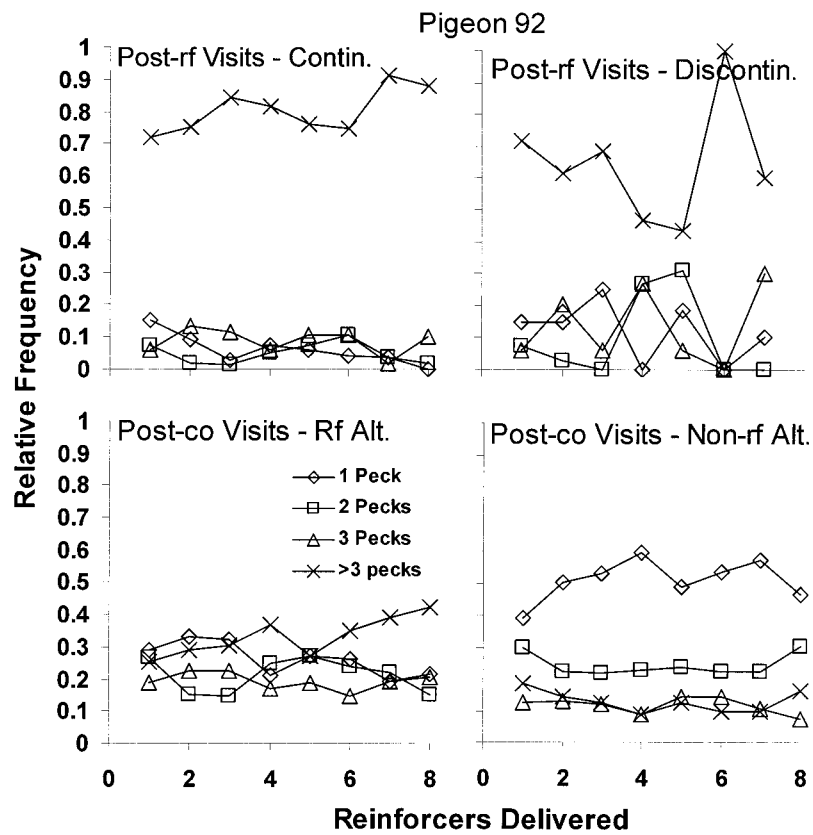


Fig. A11. Same analysis as shown in Figure 14 for Pigeon 92. See caption for Figure 14.




Fig. A12. Same analysis as shown in Figure 14 for Pigeon 93. See caption for Figure 14.

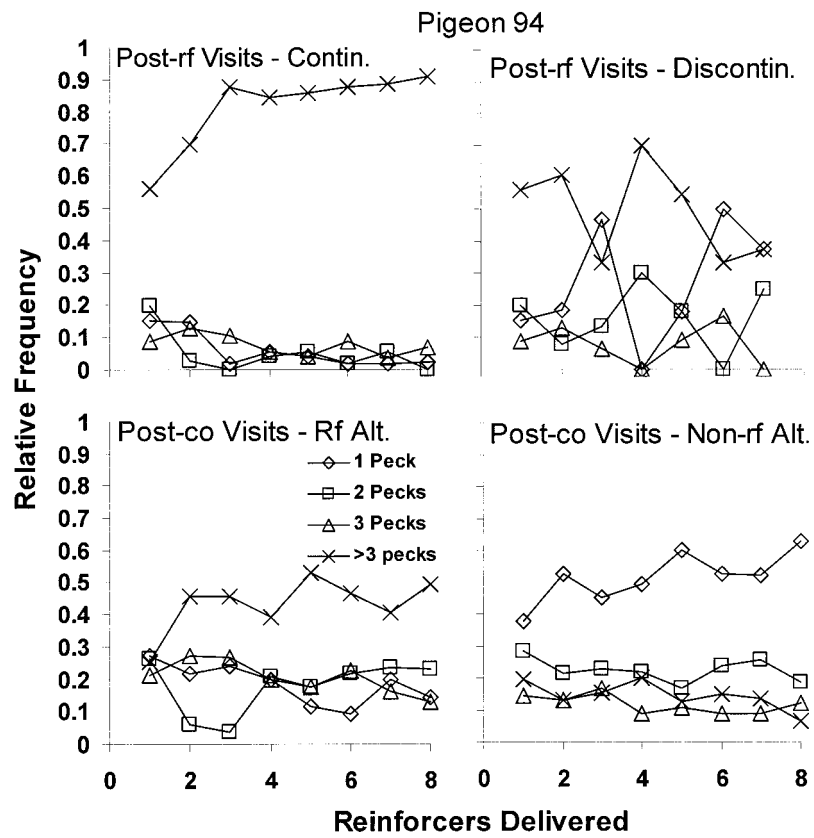


Fig. A13. Same analysis as shown in Figure 14 for Pigeon 94. See caption for Figure 14.




Fig. A14. Same analysis as shown in Figure 14 for Pigeon 95. See caption for Figure 14.




Fig. A15. Same analysis as shown in Figure 14 for Pigeon 96. See caption for Figure 14.